Health, Human Capital, and Development*

Hoyt Bleakley†

January 28, 2010

Abstract

How much does disease depress development in human capital and income around the world? I discuss a range of micro evidence, which finds that health is both human capital itself and an input to producing other forms of human capital. I use a standard model to integrate these results, and suggest a re-interpretation of much of the micro literature. I then discuss the aggregate implications of micro estimates, but note the complications in extrapolating to general equilibrium, especially because of health's effect on population size. I also review the macro evidence on this topic, which consists of either cross-country comparisons or measuring responses to health shocks. Micro estimates are 1–2 orders of magnitude smaller than the cross-country relationship, but nevertheless imply high benefit-to-cost ratios from improving certain forms of health.

Keywords: disease, income, schooling, economic growth, envelope theorem

*Draft prepared for the *Annual Review of Economics*. The author acknowledges the National Institutes of Health (P01 AG010120-14A2) and the Booth School of Business for funding support; Daron Acemoglu, Pascaline Dupas, Jon Guryan, Fabian Lange, Grant Miller, Lars Stole, and Miriam Wasserman for helpful comments; and Mariya Drachevskaya and Liang Feng for excellent research assistance. All errors and omissions are the responsibility of the author.

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

1 INTRODUCTION

Poor countries tend to be unhealthy, and unhealthy countries tend to be poor. Across the broad swath of history, improvements in income have come hand-in-hand with improvements in health. Further, the poorest countries of the world are in the tropics, an area rife with tropical disease. These stylized facts lead to a natural question: does disease hold back development? It seems possible that it would. Health is a kind of human capital as well as an input to producing other forms of human capital. Being unhealthly depresses the ability to work productively and/or the ability and incentives to invest in human capital. Taken together, these mechanisms imply that worse health implies lower income. But is this enough to account for the stylized facts above?

The correlation between health and development is hard to interpret simply as the causal effect of health on income. Third factors such as bad government or geographic disadvantages might impede both productivity and disease control. Also, health is also a normal good: when people get richer, they invest more in their own health, and exhort their governments to spend more on public health. The correlation between health and income might be one of circular and cumulative causation: health affects income and income affects health and so on.

This review is organized around the following question: in the aggregate, how much does disease depress income and human capital? The answer to this question would help us understand the global distribution of income and also inform us about the rate of return on health policies.

I start with a review of micro-empirical studies, where the evidence is most compelling, but less tightly linked to the original question, and then move to macro, where the evidence is less compelling, but better linked to the original question. I discuss evidence drawn both from developing countries today and from the historical experience of now-developed regions. The micro-empirical literature that I review is somewhat disconnected from economic theory, so I also use some stylized models to sort things out. The models help determine if a given result is of 'first order' importance, which is to say whether health has, to a first approximation, any effect at all on income.

I limit the scope of this review in a few ways. First, this review is about income, not well-being. People like being healthy, and so better health means better well-being even if the improvement in health does not increase income at all. Second, the review is not a bibliography on this topic. Instead, I discuss and critique examples of work that I group together by mechanism and methodology. Third, the topic of this review is economics, not biology. There are many fascinating detours into biological mechanisms that I could take, but the focus is instead on economic outcomes, viewed through the lens of economics. Fourth, this review considers the response of income to a given im-

provement in health, but not how to optimally design health programs. These features hopefully make this article more complement than substitute to existing reviews in this area.

Both theory and evidence suggest that we should stop thinking of health as a univariate object. Health's impact on income likely depends on how health changes (morbidity versus mortality, for example) and when (childhood, working age, or old age). Health is multifaceted, and must be treated as such.

In Section 2, I consider effects of childhood health on adult income. Early-life health could depress human capital (broadly defined) and thereby reduce lifetime income. Even if a person is perfectly healthy as an adult, damage from childhood disease may be hard to undo. Most of a person's human-capital and physiological development happens early in life. Childhood is thus a key period for human-capital building, and the burden of disease in childhood could have effects that persist throughout the life course. There may also be shorter 'critical periods' during childhood, by virtue of which some aspect of human development is hampered if it does not take place at a particular, biologically determined age.

These ideas find confirmation in four strands of the literature, all of which estimate effects of some aspect of childhood health (or input thereto) on adult outcomes. One line of this literature finds a large labor-market return to adult height, which is, to some degree, a proxy of early-life health. A second line looks at large shocks, such as famines and epidemics and shows that time in the womb is a critical period of sorts. A third set of studies shows that nutritional deprivation early in life can have long-term consequences. Finally, a group of papers examines the effect of early-life exposure to tropical diseases. But caution is indicated if applying many of these results to the question at hand, in that various studies examine a rapid shock interacted with a narrow age window, which would have different consequences than a persistent change in health.

Taking a step back, I discuss the economics of how childhood health should affect adult income. I use a simple model of human capital to integrate the variety of outcome variables used in this literature. A useful point of departure is the optimal choice of years of schooling in the standard model: a child should attend school until the marginal benefit equals the marginal cost. Childhood health plausibly affects both benefits and costs of schooling, so the model implies that the impact of childhood health on education is actually ambiguous. It is therefore difficult to interpret studies that use time in school as the only outcome.

Thinking about health and human capital in these terms leads us to the 'Envelope Theorem'. This theorem implies that improvements in health affect income by making human capital more productive, but not via more investment. This means that lifetime income would rise because child-

hood health allows you learn faster and grow up stronger. Health might also increase investments such as staying in school longer, but the discounted change in lifetime income from such investment is close zero if the marginal benefits and costs of schooling were already being equated. This analysis also affects how we evaluate evidence of complementarity across inputs to human capital, such as if early-life interventions spur later investments. The Envelope Theorem suggests that we should measure the direct effect of the intervention; whether the other inputs are substitutes or complements is of second-order importance. (This argument should be familiar to macroeconomists; it is analogous to the idea that gains in income per capita come from total-factor productivity, not from marginal changes in factors of production.)

In Section 3, I discuss the effect of adult health on adult productivity, via two distinct channels. The first channel is the direct one: being sick today reduces your ability to work today. One approach is to measure the time lost to well-defined episodes of sickness and disability. Indeed time lost to such episodes in poor countries is non-trivial, although a minor fraction of work hours. On the other hand, we know less about more subtle effects of adult sickness on income, such as reducing the quality of the labor input via diseases whose symptoms are less episodic. A second channel is more forward looking: people invest more in childhood if their human capital will not be idled by disease in adulthood. There is considerable debate in the literature on the effect of adult life expectancy on years of education and, for example, whether this effect is much greater than zero. However, this channel is of second-order importance, by the same logic as above: changing inputs that were already optimized has no effect, to a first approximation, on income.

Next, I explore the macro implications of the estimates from the micro literature in Section 4. As a point of departure, I discuss in Section 4.1 attempts to extrapolate from micro estimate to the cross-country gaps. Calculations suggest that improving health in poor countries would raise income, but these estimates range from a few percent to tens of percent, depending on the intervention studied and whether a narrow or broad measure of health was used. By these extrapolations, improving health would deliver large increases in income to unhealthy regions, although the gap between rich and poor countries is one or two orders of magnitude greater than the gains estimated from the microdata.

Incorporating the response of population, which is affected mechanically by mortality, can alter some of these results, however. Reductions in mortality brought about the majority of the increase in population in the past two centuries (Preston, 1980). Nevertheless, although lower mortality moves population onto a new growth path, where this path eventually settles depends on how fertility adjusts to the change in health. In Section 4.2, I interpret some recent studies within the

quantity/quality framework, but argue that more research on fertility is needed. Then, tying these threads together in Section 4.3, I analyze changes in human capital and population. First, I review several calibration exercises. A starting point for this exercise is to recognize that (i) mortality restrains population growth and (ii) an increase in population dilutes the per-capita supply of non-labor factors. Several studies argue that this combination can substantially attenuate—possibly even reverse—gains in income that might otherwise come from increases in human capital. This suggests that the 'Spectre of Malthus' continues today. But results based on sharply diminishing returns to labor, which are in a sense calibrated from the 20th century, may be less likely in the 21st. This is because of two large changes: urbanization and globalization, both of which make the economy less dependent on the amount of land that is locally available. I then discuss a few other issues in extrapolating from partial to general equilibrium.

In Section 5.1, I discuss evidence from cross-country comparisons in macroeconomic data. A large literature examines the relation between GDP per capita and proxies of health in a cross-section or panel of countries. The main finding is essentially the stylized fact that motivated this review in the first place. These studies are typically well done, given the methodological constraints. But the constraints are non-trivial: it is difficult to know what to make of the causality between these two variables, and the use of a single proxy variable for health muddles the policy relevance. Another difficulty in interpreting cross-sectional estimates is that we do not know the mechanism or timing with which improving health affects output.

A smaller literature, described in Section 5.2, treats the aggregate response to large health shocks. First, I discuss within-country evidence from two studies. Next, I turn to a recent paper by Acemoglu and Johnson (2007) who examine, in a cross-national panel, improvements in health technology in the 20th century. They find that decreases in mortality were followed by increases in population and concomitant decreases in GDP per capita. This suggests that returns to labor diminish rapidly, although it is hard to square this result with observed shares of labor in production. I then discuss recent criticisms of this study. I also consider explanations for the different results of these three studies, but argue that more research of this sort is needed.

In Section 6, I offer conclusions from this review. Whether we think of many of the estimates below as large or small is, in some measure, a question of perspective. 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. But from the point of view of a micro/labor

economist, some of these estimates are quite large.

2 EARLY-LIFE HEALTH AND HUMAN CAPITAL

Poor health in childhood might depress the formation of human capital. This is likely to be true for two reasons: (a) much of a person's physiological and cognitive development happens in childhood and (b) economic theory (e.g., the Ben-Porath model) shows that human-capital investments should be made early in life. Notionally,

$$\text{Early-life health} \Longrightarrow \text{Return to human capital} \Longrightarrow \text{Income},$$

which could affect lifetime income both directly and through investments. In this section, I discuss selected evidence that has been offered for this proposition. I then argue that we should focus on the direct effect, and use insights from a simple model to filter the findings in this literature.

2.1 Height as a Proxy for Health.

One indicator of early-life health is height, and numerous studies find that positive returns to being taller in developing-country microdata. Savedoff and Schultz (2000) summarize these results. One example is by Ribero and Nuñez (2000), who find that adult wages rise 7%–8% per centimeter of height in Colombia. To address omitted variable problems, they use local resources in childhood as an instrument for adult height. Whether these instruments are valid is debatable: local spending on health early in life is probably correlated with a slew of other neighborhood and parental factors. In any case, there remains the issue of how to interpret the results. Does the return to height represent the effect of childhood health on the returns to human-capital investment, the direct benefit of being more physically robust as an adult, perhaps labor-market discrimination, or some other channel?

2.2 Critical Periods.

Beyond the general importance of childhood as the base of human-capital formation, there are narrower windows that are critical for particular aspects of human development. If this development does not occur within this window, perhaps because of some input that is missing, then the individual will have cognitive or physical deficits as an adult.[1] Notionally, we can represent an

[1]There is abundant experimental evidence of critical periods in development for non-human animals. A few examples suffice: (a) mice deprived of nutrition at various points in their development exhibit a host of physiological

individual's adult human capital as $h(x_1, x_2)$, where the x_t are inputs in each time period t, and the $t = 1, 2$ are two periods of development.[2] Formally, then, $t = 1$ is a critical period if $h_1 \gg h_2$, and even more so if $h_2 \approx 0$. The key is not just that $t = 1$ is important, but also that there are no close substitutes—it is 'now or never'.

2.3 In the Womb

This section discusses evidence that the time in the womb is a critical period. A developing fetus grows, in percentage terms, at a phenomenal rate. Moreover, during the time in the womb, all of the major organs are grown from scratch.

The 'grand-daddy' of this area of research is the Dutch famine study. Towards the end of World War II, food rations were in very short supply in the Netherlands. The resulting 'hunger winter' of 1945 ended abruptly with the Allied liberation of that country. Followups found negative impacts of being in the womb during the famine on childhood growth, but minimal effects on IQ in adulthood, for example. (See Stein, Susser, Saenger, and Marolla, 1975.)

Beyond just these specific results, the Dutch famine study also provides a guide to identifying in-utero effects. Namely, if a shock is large and sharply timed, then only a narrow band of cohorts would be exposed to it while in the womb. One can assess retrospectively whether the time in-utero is particularly sensitive to the shock by analyzing outcomes across cohorts. The most affected cohorts should be visible as a 'blip' when graphing the sensitive outcome.

In recent years, other scholars have followed this model and found further evidence that the time in the womb is a sensitive period. I discuss a few here. Almond (2006) examines the 1918 influenza pandemic in a range of US datasets and finds that those born in 1919 (and therefore in the womb when the pandemic hit the US) had 2% lower income (among men) and 0.15 fewer years of schooling (both men and women). Almond and Mazumder (2008) show long-term effects on education for children in the womb during the periods of daytime fasting in the Islamic calendar. Meng and Qian (2006) find lower schooling for those born during the Great Famine in China. For the purposes of this review, results in this literature are nevertheless hard to interpret because of (i) their tendency to focus on education and (ii) the lack of comparability between these shocks

and cognitive deficits later in life; (b) monkeys that are deprived of nurturing in the first weeks of life develop behavioral problems later on; and (c) birds that are not exposed to their species' bird song during a particular window are never able to produce complex patterned song as adults. This highlights the possibility that, for humans, there might be periods in which the formation of human capital (including cognition, education, socialization, and language) is particularly sensitive to the presence or absence of certain inputs.

[2] See Heckman (2007) for a detailed modelling treatment of critical and sensitive periods.

and the baseline level of poor health and nutrition in poor countries. (I flesh out these ideas below in Sections 2.6 and 4.1, respectively.)

Several other studies examine variations in weather conditions around the time of birth. Barreca (2007) show that cohorts born in warmer and wetter years in the US South (circa 1910) had lower education, which he attributes to malaria. Maccini and Yang (2009) find effects of unusual amounts of rainfall during the year before birth on education and household income among women (but not men) in Indonesia.

Perhaps prematurely, the in-utero story seems to have become a *de facto* standard for identification of the human-capital effects of shocks. There is no reason why the time *in utero* would be the only sensitive period. For example, Glewwe and King (2001) analyze a famine in the Philipines, and find the greatest effect on two-year-old children. In addition, a study that uses a short-duration shock might not detect a longer critical period at all.

A different angle on the *in utero* question is the study of birthweight, especially among twins. For example, Behrman and Rosenzweig (2004) compare outcomes within monozygotic ("identical") female twins born in Minnesota. Comparing within identical twins removes family, genetic and temporal confounds, and they argue that differences in birthweight within the twin-pair arise because of random differences in nutrient intakes, perhaps due to different positions in utero. They find that, within the twin pair, the twin with higher birthweight had higher educational attainment and income as an adult. The external validity of this study is nevertheless open to question: are the returns to birthweight between twins the same as the returns to birthweight among singleton births? The relationship between birthweight and later-life outcomes is evidently a reduced-form one. But twin development *in utero* may be different enough from the maturation of singletons that birthweight no longer indicates the same thing.

2.4 Nutrition

Early-life health might also affect adult productivity because of nutrition, which has been studied in several interventions. (The results from famines in the previous section also were very likely about nutrition as well.)

A leading example of a nutritional intervention is the INCAP trial in Guatemala. The treatments were given to children in four Guatemalan villages four decades ago, so it has been possible to conduct a long-term followup on their effects. (Behrman, 2009, provides a compact review of the long-term followup.) The villages were split into two groups and assigned one of two nutritional

supplements: "atole" and "fresco", which were available daily to children during the trial. The atole supplement contained more protein and sugars; both had micronutrient supplements.

Results from the INCAP trial suggest quite large effects of early-childhood nutrition on adult human-capital outcomes. The findings in this paragraph refer to children exposed to atole instead of fresco in the first few years of life, using the atole-fresco difference among children exposed at older ages as a control. Male children who received the extra nutrition went on to earn wages that were one third higher than would otherwise be expected. Female children who received atole in the first few years of life attained over a year more of school. Both genders saw increases in various aptitude-test scores from the extra nutrition.

Nevertheless, a few issues with the design of the INCAP trial complicate the interpretation of these results. First, the atole was evidently more palatable than fresco, and the atole take-up was substantially higher. It is therefore hard to know whether the mechanism was more protein, or just more calories and/or micronutrients. Second, one of the control villages was drawn from a different region, and whether it represents an adequate counterfactual for the cross-cohort difference is debatable. Third, the fact that the sample is comprised of only four villages is problematic. Two control villages are probably not enough to get an estimate of the variance in cross-cohort outcomes, which is needed to make valid inferences about whether the changes across cohorts observed in the treatment villages were unusually large. Because these effects are so large, and the policy relevance so clear, it would be useful to obtain further confirmation of this specific result for nutrition supplementation and income, perhaps in a well-designed retrospective study using the rollout of nutrition supplements that followed on the heels of the INCAP trial.

Finally, the importance of iodine in human-capital development has been extensively studied. The iodization of salt is very inexpensive, and so should be regarded as an easy intervention. Iodine deficiency has been related to mental retardation and intelligence in observational studies. Further, a meta-analysis (DeLange and West, 2003) of several interventions finds that remedying iodine deficiency early in utero raises the child's IQ by 5 points, or one half of a standard deviation.[3] None of these studies has results for adult income, however.

2.5 Tropical Parasites

The childhood morbidity coming from infection by tropical parasites may substantially depress human capital and subsequent income. I focus mostly on malaria and hookworm, parasites that

[3] Policy documents and editorials often make reference to a 10-15 point increase in IQ from iodine supplementation, but this number appears to be chosen from the study with the largest point estimate.

are generally found in tropical regions. Both diseases have potentially severe chronic symptoms, notably anemia. Malaria, whose acute symptoms include fever, headache, and nausea, can result in death, but the case-fatality rate is actually low, especially among the strains predominant outside of subsaharan Africa. Hookworm is rarely fatal. Among children, hookworm and malaria have been associated with stunting of physical and cognitive development. It is not a leap to imagine that this reduces the return to human-capital investment. While age brings partial immunity, the damage from childhood exposure to these parasites may be hard to undo.

Research into the human-capital consequences of parasite infection during childhood falls into two broad categories: observational and experimental. Each type of research has its own advantages, but the two groups are, generally speaking, complements rather than susbtitutes. Experimental studies allow us to be well assured of the causality, but are typically limited to the short term. Observational studies, on the other hand, can consider longer term outcomes if they are based on the retrospective analysis of health conditions. Thus many of the observational studies that I review here consider eradication campaigns that took place some time ago.[4]

While one can never rule out that results from an observational study are driven by some unmodeled factor, the commonalities across the campaigns studied in the literature suggest a 'recipe' for identification. First, the studied campaigns brought about large and sudden changes in the disease environment. Second, the countries considered have (a) less-infected areas that can serve as a comparison group, and (b) microdata available that cover the relevant sets of cohorts. And third, these campaigns began because of critical advances in health technology and spending that originated outside the affected regions, which should mitigate concerns about reverse causality.

One example of this line of research is a study (Bleakley, 2007a) in which I analyzed the impact of a hookworm-eradication campaign[5] in the southern US, where in 1910 children had infection rates of 30%–40%. Children growing up in areas that benefitted from the anti-hookworm campaign saw large increases in literacy, school attendance, and income, relative to earlier cohorts. The changes in these outcomes coincide with childhood exposure to the eradication efforts, which further suggests that we are measuring the effect of deworming rather than some subregional trend. According to my estimates, childhood exposure to hookworm depressed adult income significantly in the South.

[4]In the 20th century, malaria and hookworm were widespread even outside of tropical regions. These parasites were, in many cases, the focus of 'big push' campaigns for eradication.

[5]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.

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 your adult income by 43%. (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.

Experimental studies also find positive impacts of deworming on human capital, albeit at shorter horizons. Two examples are Miguel and Kremer (2004) and Bobonis, Miguel, and Sharma (2006). Both papers find that intestinal worms reduce school attendance. (Miguel and Kremer did not, however, obtain strong results for achievement tests.) An important feature of the Miguel/Kremer paper is the assignment of treatment at both the individual and school level. Disease transmission from control to treatment group will tend to reduce treatment/control differences, which they show.[6] (They interpret some of their effect on attendance as working through non-health spillovers, although it is possible that measurement or specification error on the health side could generate something resembling a non-health spillover.)

In another study (Bleakley, 2007b), I 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), and in Brazil, Colombia and Mexico (circa 1955). The basic finding of the study is that cohorts born after eradication had higher income and literacy as adults than the preceding generation, relative to cohorts from non-malarious areas. The change across cohorts coincides with childhood exposure to the campaigns. The magnitude of the change in income associated with childhood malaria is substantial. Reduced-form effects on income, when comparing the least malarious to the most malarious areas within a country, are on the range of 12% (in the U.S.) to 40% (in Latin America). Although it is impossible to completely rule out that the intervention had effects through channels besides estimated malaria infection, the results that suggest that persistent childhood malaria infection reduces adult income around 50%. The normalized effects are similar across the four countries.

Two other papers find that childhood exposure to malaria-eradication campaigns had impacts on human capital. Lucas (2009) shows that women born after malaria eradication in Sri Lanka and Paraguay completed more years of schooling and were more literate. Her estimates for literacy are similar in magnitude to those by Bleakley (2007b). Cutler, Fung, Kremer, Singhal, and Vogl (2009) analyze the malaria-eradication campaign in India and find malaria exposure in the first few years of life had little effect on education, but raised household expenditures (presumably because

[6] See also Bundy, Kremer, Bleakley, Jukes, and Miguel, 2009; and Philipson, 2000.

of higher productivity) among males. Their headline number for expenditures is lower that what I reported in the previous paragraph, but estimates are similar once adjusted for their focus on just the first years of life (instead of all of childhood) and for their sample being 20-25 years old, ages at which the returns to skill are lower.

2.6 The Role of Schooling, and the Envelope Theorem

The outcome variables above were a mix of income and education. How do we combine these results into a coherent whole? First, recall that years of schooling is an input to the production of human capital; it is not synonymous with human capital itself. This fact affects the interpretation of the literature above.

Consider a textbook model in which the individual chooses time spent in school so as to maximize lifetime income. Let the benefits of schooling be $b(e, h)$ and the costs be $c(e, h)$, where e is years of education[7] and h is health (thought of here as lower childhood morbidity or better developed cognition). The benefits b include the appropriately discounted sum of future earnings. The costs c are both direct and opportunity costs of schooling. The usual assumptions[8] are that the marginal benefit of schooling declines with more time in school, but that the marginal cost rises: $c_{ee} > 0$ and $b_{ee} < 0$, where subscripts denote partial derivatives.

A child should stay in school as long at the marginal benefit exceeds the marginal cost. At an interior solution for schooling he will equate marginal benefits and marginal costs, $b_e = c_e$, which maximizes lifetime income with respect to years of schooling. If you are even close to the optimum, marginal benefit is close to marginal cost, and you are pretty much indifferent between attending more or dropping out. Thus, compelling, tricking, or bribing people to spend more years in school has its effect when the net gains from schooling have been largely exhausted.[9] (Numerous fashionable policy initiatives, such as conditional cash transfers, do this to some extent.)

[7] More precisely, think of e as the time when you leave school and enter the labor market, as in the Ben-Porath model. I will discuss inframarginal time in school (such as attendance) below.

[8] These assumptions are intuitively appealing. At low levels of schooling, the marginal benefit is high: you are learning skills—the alphabet, for example—that your brain can digest and that have a high payoff in society. At the same time, the marginal cost—specifically, the opportunity cost of foregone earnings—is low: six-year-old kids are not very productive in the labor market. The early years of primary school, by this calculation, have a much higher benefit than cost. As you grow up, however, your potential wage grows, and thus the marginal cost of schooling rises. All the while, the marginal benefit declines with more time in school, your brain being less adept at learning the now-more-esoteric subjects.

[9] Note that (i) the net income gains from schooling are benefits net of both direct and opportunity costs, and (ii) everything is properly discounted into present-value units. The reader might argue (normatively) that an individual could discount the future too much, but the relevant calculation here is about the welfare of that individual. Therefore the correct discount rate to use is the individual's. Note further that gains from schooling here are different from the Mincerian returns to schooling, which are the income/education gradient at a point later in life.

Perhaps counterintuitively, the effect of childhood morbidity on years of schooling is actually ambiguous. The optimal response of schooling to health becomes $\frac{de^*}{dh} = -\left[(b_{ee} - c_{ee})/(b_{eh} - c_{eh})\right]$. By assumption, the numerator is negative. The working assumption of the literature reviewed above is that childhood health increases the marginal benefit of schooling: $b_{eh} > 0$. But it is probably also the case that $c_{eh} > 0$: a healthier child can earn more in the labor market. This term makes the denominator of ambiguous sign. In words, reducing childhood morbidity raises years of education if and only if childhood morbidity depresses the marginal benefits of education *more* than it reduces the marginal costs (including opportunity costs, such as foregone wages). This result is true regardless of the sign of b_h.

With this simple and entirely standard model, we see that years of schooling (e) is not a sufficient statistic for measuring the impact of early-life health on lifetime income. It is possible that when health improves, lifetime income goes up, but years of school declines. All that needs to happen is that being a healthy kid raises the wage more than the returns to the schooling. Indeed, any combination of sign($\frac{de^*}{dh}$) and sign(b_h) is possible. Herein lies the problem for using education as a sufficient statistic. Yet, many of the studies reviewed above measured responses of years of schooling, but not income.[10]

At this point, some readers might question if schooling being chosen to optimize income is a good assumption, although this does not rescue the view that years of schooling is the central outcome. First, perhaps people maximize lifetime utility rather than lifetime income. Fine, so the analysis above is only about welfare. Second, some would argue that the rise in opportunity cost from improved health is low because of unemployment, but note that unemployment also reduces the marginal benefit of schooling. Third, suppose that there is some constraint that prevents children from going to school: for example, a credit-market imperfection. In this case, although more education would increase lifetime income, we do not know whether improving childhood health would unambiguously predict going to school more. It seems just as likely that health increases the opportunity cost without resolving the problem in the credit market.

Next, the quantity of schooling is not an important channel through which childhood health impacts income. The reason, simply stated, is the Envelope Theorem.[11] Consider the individual's

[10] One might argue that the preponderance of positive results for education on health would suggest either that my analysis here is wrong or that returns to education always increase more than child wages when childhood health improves. But I am inclined to wonder how many negative results for education wound up on the cutting room floor.

[11] The Envelope Theorem is a fundamental result in microeconomic theory that describes how the value of an optimized function changes in response to a small change in the parameters of the problem. At first blush, this might seem like a complicated task because all of the input choices could change in response to the new parameter vector. However, the inputs were chosen optimally before, so that adjusting them has no first-order effect on the optimized

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^* w.r.t. h:

$$\frac{dy^*}{dh} = \left.\frac{\partial y}{\partial e}\right|_{e^*} \times \frac{de^*}{dh} + \left.\frac{\partial y}{\partial h}\right|_{e^*}. \tag{1}$$

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

This result seems at odds with the way education treated in some empirical microeconomic research. In my experience, education is often presumed to be a central channel between early-life events and income. Indeed, it seems to be used as a plausibility check by some, such that a result for income cannot be correct if it is not 'explained' in large measure by years of schooling. This is ironic: instead of it being a plausibility check, it is an irrelevancy check. Limiting ourselves to studies of health and income in which education is the main mechanism is a recipe for minimizing our estimates of the effect of health.

These results matter for the interpretation of the studies reviewed above. In the interest of even-handedness, let me pick on some of my own research first. One of the better documented critical periods from psychology is associated with language acquisition. In Bleakley and Chin (2004), we examine childhood immigrants to the US who arrive before and after this critical window, and use migrants from English-speaking countries as a comparison group. We find that arriving after the critical period earn as much as 15% lower wages as adults, and attend 1.5 fewer years of school. One approach would be to say that this is an enormous effect of language because we found a large increase in earnings and in schooling. But those extra years in school were not free; they were instead purchases by not working during that time. At rates of return to schooling for that sample, the increase in schooling was only slightly more than compensated by the increase in wages later on. It makes a big difference whether we are talking about an effect on income that is 15% (the unadjusted) or an effect that is closer to 3% (adjusted for schooling).

How does this insight apply to the literature on health and human capital above? First, all of the literature above whose only outcome variable was years of school are, by the envelope logic,

value of the function. There are also interaction effects between the change in the parameters and the change in the inputs, but these are of second-order importance.

measuring changes that are not of first-order importance to lifetime income. Next, consider some examples that treat both education and income. On the one hand, Almond (2006) measures decreases in earnings around 2% for men born in 1919, but at least half of this is due to having spent less time in school. Indeed, for a borrowing interest rate of 13% or more, the 1919-born men actually had higher lifetime income. Among women, Almond estimates that being born in 1919 reduced education, but with no concomitant effect on earnings. While it is difficult to estimate the returns to schooling if much of that sample was in home production, spending more time in school and getting no return in the labor market (*ceteris paribus*) can hardly be viewed as a good thing that the women born in 1919 missed out on. (Both genders suffered from more health trouble if they were born in 1919, which I presume reduced their utility. None of the preceding analysis is meant to minimize this.) A different example is by Maccini and Yang (2009), who find increases in income proxies for women in Indonesia who were born after favorable weather conditions, and they argue that these results are entirely explained by increases in education. Filtering their results through the standard theory of schooling suggests no sensitivity at all of lifetime income to early-life rainfall. (They also report zero effects for males in their sample, which does not require as many logical steps to interpret as a zero effect on income.)

So how might childhood health increase human capital and, thereby, income? The model provides further insight. The second term of equation 1 measures the direct effect of health on labor productivity, evaluated at e^*. It is instructive to decompose this direct effect into two parts:

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

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.

Several studies reviewed above provide examples of improved health increasing income, but not principally via marginal years of schooling. In the INCAP trial (Behrman, 2009), males exposed to the extra nutrition during their first three years experienced large increases in income, but very little change in time in school. Estimates from the campaigns against tropical disease (Bleakley; 2006b, 2007a; Cutler *et al*, 2009) also illustrate the problem with using the quantity of education as the chief channel of childhood-morbidity effects on income. In no case does schooling account

for more than a quarter of the effect of productivity. Indeed, in one case (Mexico), adult income rises with less childhood exposure to malaria, but years of school actually decrease.

Beside measuring the effect on income with income data, how else can we detect first-order effects of childhood health on income? Outputs like literacy and achievement/IQ tests seem like good measures as well, although the eventual benefit from acquiring such skills would need to be measured and weighed against the extra time in school (if any) that was required to learn them. One can also look at the intensive margin: more frequent attendance[12] or school progression, for example. Especially for young ages, these measures are clearly inframarginal; an example is Bobonis, Miguel, and Sharma (2006) who analyze anemia and pre-school attendance. Nevertheless, for the purposes of this review, we would need to measure the labor-market return to these activities, if we want to know the resulting increase in income.

2.7 Critical Thoughts on Critical Periods

Blip-blip under the lamppost. Many of the studies use what I term a 'blip blip' design. They analyze a sudden shock (blip #1) interacted with a short critical period (blip #2). The compact nature of these two blips are what help to identify the estimation. Getting the research design right is clearly a good thing in isolation, but we should take stock of the phenomena that we ignore as a result. In this case, if you want to identify an effect that has a longer window of sensitivity or with a shock that is less compactly timed, you might be out of luck. In my own work on tropical disease, I attempt to measure the effects of exposure during all of childhood using campaigns that took effect over five or more years, but it required a combination of long spans of data (75+ years of cohorts) and diseases that likely had large impacts, a 'perfect storm' that might not be repeated for other questions.

When thinking about how much poor health impedes development, we should contemplate a large and persistent reduction in disease. But—however well identified—short-term variation in [insert disease here] is likely to have economic impacts that differ from those due to long-term changes. In terms of the model of critical periods above, a blip-blip design measures the differential effect of early exposure, or $h_1 - h_2$, but, for the question that motivates this review, we want to know the level of both h_1 and h_2.

[12] One doubt is that perhaps students with significant absenteeism would not be allowed to advance in school, and thus these measures will show up in years of schooling. This may be true to some extent, but in an area with heavy childhood disease burdens, it seems more likely that the curriculum would be diluted to accomodate this constraint.

The Earlier, The Better? A recent theme in policy circles has been the importance of intervening early in childhood. The above discussion on critical periods would seem to substantiate this. But we should exercise caution in interpreting sensitivity to early inputs as evidence of critical periods, or even as evidence that early intervention is warranted.

Here I present a model that exhibits excess early sensitivity, but does not have a critical period, nor are the correctly measured returns to early intervention actually higher. I return to the model of a critical period shown above, but strip out the features that yielded a critical period. Lifetime income is $h(x_1, x_2)$, where the x_t are investments made in periods $t = 1, 2$, corresponding to early and late childhood. Suppose that for any $x_1 = x_2$, the marginal products are the same for each x; that is, $h_1 = h_2$, where subscripts on the h refer to partial derivative w.r.t. the numbered argument. This implies that there is no critical period in x. Suppose further that the x are complements,[13] so that $h_{12} > 0$.

What are the returns to exogenous changes in the inputs in this model? Suppose that the additional input was unanticipated, as would likely be the case in trial intervention. The response to additional input in period 2 is $\frac{dh}{dx_2} = h_2$, while the response to additional input in period 1 is $\frac{dh}{dx_1} = h_1 + h_2 \frac{dx_2}{dx_1}$. Because the decision about x_2 is made after x_1, the optimal x_2 can incorporate the new information about x_1, whereas the opposite is not the case. The x are complements, so $\frac{dx_2}{dx_1} > 0$. It follows that $\frac{dh}{dx_1} \geq \frac{dh}{dx_2}$. The response to earlier inputs is larger than for later inputs, because you can invest more in period 2 in response to the intervention in period 1. This result is completely mechanical and does not depend on there being critical periods or even excess sensitivity to early inputs in the production function. Instead, the result depends only on the fact that early-childhood-investment decisions come before those in late childhood, which does not seem like a controversial assumption.

What empirical findings would come from a study of interventions in x_1 or x_2? If the x_t are chosen to maximize lifetime income h, then $h_t = c$, for $t = 1, 2$ and c being the unit cost of the inputs x_t. It follows that

$$\frac{dh}{dx_2} = c \text{ and } \frac{dh}{dx_1} = c\left(1 + \frac{dx_2}{dx_1}\right) > c.$$

In words, the increase in lifetime income is greater if the child is given a little extra input in period 1 instead of period 2. Perhaps the researcher who obtains this result would stop here, but is this really evidence that 'the earlier, the better?' No, because the increases in input costs exactly match the increases in income. (I could go through the case in which x_2 is a time input like years

[13] Cunha and Heckman (2008) refer to this sort of complementarity as "dynamic complementarity".

of schooling, and–even worse–if the researcher used x_2 as the main dependent variable instead of income, but this is treated above.)

A remaining concern with this analysis, however, is that early-life investments might be chosen suboptimally, perhaps by parents who do not fully account for their children's welfare in their decisions. Inputs to human capital might not be set optimally, but this issue cannot be resolved by demonstrating complementarity with other inputs. Instead, what is required is a comparison of gross returns (the $\frac{dh}{dx_t}$) with an accounting of all associated costs, including those downstream.[14] While the conventional wisdom among economists is that 'crowd out' is bad, this does not imply that 'crowd in' is a free lunch. The resources that are 'crowded in' themselves have opportunity costs, and again the Envelope Theorem is relevant.

The insights from the model above should guide our interpretation of observational studies as well. An amusing example is the much discussed Stanford Marshmallow Experiment (Shoda, Mischel, and Peake, 1990). In that study, young children are asked to refrain for 15 minutes from eating a marshmallow that had been placed temptingly in front of them, and the prize for waiting was a second marshmallow. The children who were able to delay gratification were also found to score significantly better on standardized tests a decade later. No doubt they will do better in the labor market as well. But is there a large return to this skill? Perhaps, although presumably much of the later increase in human capital comes from working hard, which is not without cost. Empirically, whether we can measure the fundamental return to such skills depends on whether we can cost out their correlated investments.

3 ADULT HEALTH AND INCOME

Adult health might also affect income, through two human-capital channels.

1. Adult health $\Longrightarrow$ Income. That is to say: directly. This is most easily understood as illness today reducing the ability to work today.

2. Adult health $\Longrightarrow$ Return to human capital $\Longrightarrow$ Income. The return here, unlike above, is akin to a lower depreciation rate: the utilization rate of human capital can be higher if it spends less time idled because of disease. If this is expected, early-life investments in human capital should increase.

[14] For example, Cunha and Heckman (2008) estimate a selection model, which uses the idea of revealed preference to identify unobserved costs and benefits. It seems unlikely, however, that a purely reduced-form strategy would fully account for all of the costs of induced or correlated investments.

3.1 Contemporaneous Effects of Health on Income

Perhaps the most obvious channel through which health reduces income is that you work less when you are sick. But the pertinent question is by how much. Researchers have addressed this question with several strategies.

The first approach for measuring the magnitude of this channel is by accounting for lost output during episodes of acute sickness. It is understandable that researchers would focus on acute episodes because the causality is clear, at least in the Granger sense. However, not all diseases can be well characterized simply as a sharply defined episode of acute morbidity. Ashraf, Lester, and Weil (2009) provide a thorough accounting of such effects across a broad set of diseases. Relative to rich countries, the additional labor input lost to well identified episodes of sickness in poor countries is non-trivial, but not a number large enough to close much of the gap between these two groups. At some level, this fact should not be surprising. Swedish people work, if anything, less over the life cycle than Swazilanders, in spite of Sweden being less disease ridden.

A second approach consists of observational and intervention studies—that is, either in a cross-section or in a panel following up on treatment—of income and specific diseases. An early example is by Schapiro (1919), who examines the productivity increases among plantation laborers in Costa Rica following deworming treatments. Almost all the workers were infected with hookworm at baseline, compared with approximately 10% after treatment. He found that productivity six months after treatment was 15%-30% higher, and the area cultivated had expanded by 50% the following year. Weisbrod *et al.* (1973), on the other hand, found small effects of various parasitic diseases on earnings in a cross section of laborers on Saint Lucia. (The case of plantations is noteworthy: large land-owners might have better incentives than individual farmer to manage the transmission of such diseases.) Relatedly, Brinkley (1997) attributes much of the decline of output in the US South after the Civil War to the increase in tropical disease.

A third method is the experimental evaluation of treatments for specific diseases. In spite of the econometric advantages in conducting these studies, these sorts of studies—specifically ones that measure impacts on contemporaneous income—are few and far between. A leading example in this vein is by Thomas *et al.* (2003), who conduct a randomized trial of iron supplements in Indonesia. They find increases in contemporaneous adult productivity, and argue that the main mechanism is the reduction of anemia. A feature of this study relative to earlier work is that it follows up individuals rather than basing the sample on particular workplace, where selection effects could bias results and where returns to health might be muted since occupational choice is limited.

3.2 Adult Health and the Return to Schooling

A decrease in working-age mortality extends the horizon over which human-capital investments can be utilized. Recall the model above, but now think of h as adult life expectancy. A decline in mortality means the asset called 'human capital' now depreciates more slowly, which increases the benefit of going to school: $b_h > 0$. Even with schooling held fixed, this generates a first-order gain: people can spend more time working because they do not die as quickly. This has nothing to do with education; it is simply a change in expected value stemming from the change in discounting. Nevertheless, the reduction in working-age mortality also gives people the incentive to spend more time in school, unambiguously. Unlike the change in h above, the reduction in adult mortality should not affect the opportunity cost of going to school as a child, so $c_h = 0$ and therefore $c_{eh} = 0$. By standard assumptions, $\frac{de^*}{dh} = -(b_{ee} - c_{ee})/b_{eh} > 0$.

However, the relevance of this theoretical result for education has been the subject of debate. Preston (1980) argues that this effect is small relative to secular increases in education. In contrast, Kalemli-Ozcan, Ryder, and Weil (2000) present a calibrated model arguing that time in school should respond to a large degree. Nevertheless, Hazan (2006) shows that, historically, while adult life expectancy has risen in developed countries, hours worked have actually declined, and thus the effective discount rate that should be applied to human-capital investments has gone in the opposite direction from the horizon story. Conversely, Jayachandran and Lleras Muney (2009) examine declines in maternal mortality in Sri Lanka and estimate that time in school rises 0.11 year per one year of additional expected adult life. Fortson (2009) considers schooling and AIDS in Africa, and she estimates of 0.386 year of schooling per year of adult life expectancy. These results are probably upper bounds, given that the health changes might affect schooling via other aspects of household production.

This debate notwithstanding, this mechanism is not of first-order importance for income in a standard model. Higher schooling delivers a marginal benefit in the form of higher wages in the future, but comes at the cost of spending less time working early in life. Evaluated near the old optimum, this marginal benefit is still very close to the marginal cost. The horizon channel of education is therefore only of second-order importance for lifetime income. (In Harbergerian terms, the benefit of more education here is a triangle rather than a rectangle.) The action for expected income works through the change in health, not through the resulting changes in inputs. Readers will recognize the Envelope Theorem at work, yet again.

4 GOING FROM MICRO TO MACRO

4.1 Extrapolation

The results above suggest that health conditions can affect productivity. But how do they help us gauge health's contribution to the gap between rich and poor regions? To begin, consider an extrapolation of the following form:

$$\begin{array}{c}\text{gain in } \log_n(\text{income}) \\ \text{per unit increase in} \\ \text{health}\end{array} \times \begin{array}{c}\text{gap in health} \\ \text{between rich and} \\ \text{poor countries}\end{array} = \begin{array}{c}\text{gain in } \log_n \text{ income in poor} \\ \text{country from raising health} \\ \text{up to rich country's level}\end{array}$$

Combining estimates from well-identified micro studies with survey data on health, we should have the ingredients necessary to make this calculation. Readers may balk at this exercise as it is a somewhat wild-eyed extrapolation. Granted, but the point of the exercise is to define an order of magnitude, not an exact number.[15]

One approach would be to measure the effect of health on years of schooling and extrapolate from there, multiplying by standard (i.e., Mincerian) returns to schooling. This would be wrong, for reasons that were detailed in section 2.6. Again, if education is chosen close to optimally, the first-order effect of increasing education on lifetime income is much smaller than the Mincerian return. Instead, we need to measure the impact of health on income somehow, otherwise we are missing that first term in the equation above.

Income was indeed an outcome variable for a few studies related to epidemics or famines, but this is not enough to start extrapolating. The basic problem is that we do not know how to normalize these results into units that are informative about the gap between poor and rich regions. An example might clarify this issue: consider a cohort born in a rich country, but in a year with some bad shock. By analogy, the cohort most exposed to the bad health shock is like someone born in a poor country's health environment. But is this shock like the environment in poor countries by a lot or just a little? We do not know, and so we cannot use these results to extrapolate. Instead, we need to consider types of health that are reasonably well defined and arguably comparable across place and time, otherwise we are missing that second term in the equation above.

[15] It is worth noting also that this is a *linear* extrapolation, which discounts the possibility of a 'poverty trap'. An artfully placed nonconvexity in the causal response of income to health, $y(h)$, might generate large differences between the macro response and the typical micro estimate, the latter of which is based upon relatively small changes in health. However, the discussion above should illustrate that it is hard enough to estimate the first derivative of $y(h)$. A poverty-trap model rests instead on assumptions about the third derivative of $y(h)$; perhaps we are in a situation where theory is ahead of measurement.

A notable attempt at such an extrapolation for a broad measure of health is by Weil (2007). Weil uses some of the micro studies on returns to height (cited above) as a benchmark. Because of data problems, however, he has to relate these results to the cross-country question by translating the units from height into life expectancy, by way of the adult survival ratio. The two additional reduced-form 'translation' terms in Weil's calculation add error to the extrapolation, but this exercise might still be useful for obtaining an order-of-magnitude estimate.

The results from Weil's extrapolation suggest a large, though not pre-dominant role for health. In terms of the variance in output per capita across countries, health explains a fraction that is comparably sized to what is explained by either physical or human capital. Further, these three factors combined explain a narrow majority of the variance in output per capita across countries. Nevertheless, a substantial minority of the variance is still unexplained by measured factors.

An alternative way to extrapolate from the micro evidence is to consider narrower measures of health, rather than attempting the calculation for single proxy of general health. If the measure of health is tightly defined, it might be easier to translate estimates from one context to another, and thereby avoid the 'units' problems in Weil's exercise. But this has the disadvantage that the measure of health is not a comprehensive one.

An example of this approach is by Behrman and Rosenzweig (2004), who extrapolate their estimates of the effects of birthweight to cross-national data. In their study, they use within-twins differences in birthweight and adult income from a sample in Minnesota. They report, using their within-twin estimates, that birthweight differences explain around 1% of the variance in output/capita across countries. (However, about one third of their estimates for income worked through years of schooling, which suggests that their extrapolation should be a third lower.) On the other hand, in a bivariate regression on cross-national data, average birthweight explains around 45% of the variance in GDP/capita. These two results suggest that most of the birthweight/income relationship across countries is due to reverse causality and omitted third factors, rather than the causal effect of health. However, their within-twin estimates from Minnesota may not extrapolate well to poor countries, and this depends on the relative size of selection and scarring effects across areas.

The evidence from historical campaigns against malaria and hookworm suggest that eradication would raise the income of tropical countries. I extrapolate based on available infection data and my estimates discussed above. If hookworm were eradicated today, income would rise in northern Brazil and wet/tropical Africa by 11% and 24%, respectively. If malaria had been eradicated in the 1950s in the tropical regions, northern Brazil would have seen a 35% gain, and the wet tropics

of Africa would have seen an increase in income of approximately 40%. Note that this is via the mechanism of childhood exposure only.

4.2 Whither Population?

A potentially large effect of improving health is to change the size and composition of the population. This effect might even be more important than the human-capital effects of health on development. Increases in population come from people not dying as quickly as before, and from changes in fertility and fecundity that follow improvements in health. To help sort out the contributions of each, let me start by appealing to a higher authority on the historical experience:

> Not a shred of doubt remains that the vast majority of the acceleration in world population growth during the 20th century is attributable to mortality decline rather than to a rise in fertility. Preston (1980, pp315-16.)

Nevertheless, for a permanent reduction in mortality, the eventual deceleration is because of fertility, and the magnitude of the shift from the previous baseline for population depends on how fertility adjusts. So, we have to consider fertility, too.

Unfortunately, our understanding of fertility's adjustment to large changes in health (among other things) leaves something to be desired. One common approach in demography is to allow for an *ad hoc*, mechanical adjustment of fertility to a rate that stabilizes population growth at the previous level, but this seems a poor description of the data (Preston, 1980). Furthermore, even within this framework, the long-run change in population depends crucially on the speed of fertility adjustment (Ashraf, Lester, and Weil, 2009). The demographic transition (or lack thereof) is marked by substantial heterogeneity, both historically in the developed world and recently in less-developed countries. This sort of instability in the reduced-form relationship highlights the complicated nature of the comovement between health and fertility.

The contribution of economics is often to introduce prices into thinking about human behavior, and fertility is hopefully a good example of this. Consider the standard quantity-quality (q^2) model, which has two features. First, think of child rearing as having distinct two outputs: the number of children (quantity, or n) and the average human capital of children (quality, or q). Parents will demand quantity and quality according to their prices: p_n and p_q, respectively. Second, by some mechanism (perhaps the budget constraint), quantity and quality are substitutes. The first point finds confirmation in several recent empirical studies of fertility. Acemoglu and Johnson (2007) examine the large declines in mortality (especially infant mortality, so $p_n \downarrow$) that took place

circa 1950, and find large, positive responses in birth rates. When analyzing declines in maternal mortality ($p_n \downarrow$) in Sri Lanka, Jayachandran and Lleras Muney (2009) find an increase in fertility. But does fertility respond to p_q? Fabian Lange and I (Bleakley and Lange, 2009) examine hookworm eradication in the southern US. This change reduced p_q in that it was now easier to rear an educated child, while p_n was essentially unaffected because the disease was rarely lethal. Fertility declines with the reduction in hookworm, which is consistent with the q^2 model.[16] Recent work on malaria fits in this framework, but suggests that a generational perspective is useful as well. Adrienne Lucas (2009) 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 non-malarious parts of Sri Lanka. 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, and thus the opportunity cost of mom's time was higher. Even if the q^2 interpretation is not correct, however, these results are hard to square with a univariate model of health.

There are also fertility shifters that themselves might interact with health. Schultz (1985) showed that fertility went down when female-intensive crops became more lucrative, thus increasing the opportunity cost of female time. Urbanization is another phenomenon that affects both health and fertility. Contraception programs can also affect both fertility and health (Miller, 2005; Joshi and Schultz, 2007). How population responds to health depends on the presence of these factors.

4.3 Beyond Partial Equilibrium

When the health environment changes, so do the size and composition of the population. These resulting shifts in quantities supplied should all change relative prices. Accordingly, we need to think in general equilibrium. Below, I discuss some work in this vein and highlight mechanisms to bear in mind when we go beyond partial equilibrium.

[16] The relative magnitudes of quantity and quality responses estimated by Bleakley and Lange are similar to the aggregate comovements of these variables, which provides a quantitative confirmation of the model as well. However, studies using twin births, such as the seminal article by Rosenzweig and Wolpin (1980), tend to find smaller responses of quality to quantity. (See also Angrist, Lavy, and Schlosser (2005), Black, Devereux, and Salvanes (2005), Caceres (2004), and Qian (2006).) It is nevertheless an open question whether rationing parents away from their optimal fertility choice (by twinning, e.g.) identifies the relevant mechanism in the q^2 model. A preferred test of the model is to examine responses of q and n to changes in prices, an exercise that maps more closely on to the historical changes. Put differently, if the goal is to understand the historical experience, then we should recognize that secular trends did not arise because parents were forced to have fewer children, but rather because they adjusted their fertility decisions to changing prices.

Calibration Exercises. Several studies consider the quantitative impact of health on output per capita in a calibrated, general-equilibrium framework. I discuss some examples here. Barlow (1967) considers the impact of malaria eradication, although that study does not consider the human-capital aspect quantitatively. Young (2005) gauges the effect of HIV/AIDS. A paper with a more comprehensive coverage of diseases is by Ashraf, Lester, and Weil (2009), who use a standard neoclassical framework. They combine a demographically driven change in effective labor with both a fixed factor and a slowly adjusting capital stock. They analyze a single-sector economy, and fertility is modeled as a partial adjustment to stable growth, as described above. All three studies present a fairly pessimistic view of the effect of health improvement on output per capita. Central to these results is an assumption of diminishing returns to labor, which arise if some other productive factor, such as land, is in fixed supply.

The effects of land dilution can indeed be large in a model land-dependent economy. Consider an economy characterized by a constant-returns-to-scale (CRS) production function: $Y = h(K, L, F)$. The inputs to h are capital (K), labor (L), and a fixed factor (F). The reproducible factors, K and L, can adjust, while the fixed factor is assumed to be invariant. The response of income per capita to an increase in labor is, in elasticities, as follows:

$$\frac{d\ln(Y/L)}{d\ln L} = \left(\frac{\partial \ln h}{\partial \ln L} + \frac{\partial \ln Y}{\partial \ln K} \times \frac{d\ln K}{d\ln L} - 1\right). \tag{3}$$

The first term reflects the marginal product of labor. The second term comes from the marginal product of capital $\times$ how capital adjusts to the higher population. The third term is unity: income per capita declines mechanically when there are more capitas. Because the production function h is decreasing returns to scale (DRS) in reproducible factors, this derivative is less than one.[17] But what is the magnitude of this effect? Start, for purposes of discussion, with the assumption that h is Cobb-Douglas: $h = K^{\alpha}L^{\beta}F^{1-\alpha-\beta}$. This implies that the elasticity of income per capita to labor is $\beta + \alpha\frac{d\ln K}{d\ln L} - 1$, which is on the interval $[\beta - 1\ ,\ \beta + \alpha - 1)$. A typical developing country has a natural-resource share of between 10% and 20% (Weil and Wilde, 2009, Figures 1, 2). If the fixed-factor share is 15% and labor has a two-thirds share of the reproducible factors, then the bounds on $\frac{d\ln(Y/L)}{d\ln L}$ are $[-0.43, -0.15)$. In words, an increase of the labor force of 10% would depress output per capita by at least 1.5%, and possibly as much as 4.3% if physical capital does not adjust. These bounds would be $[-0.53, -0.3)$ for an economy with a fixed-factor share of 0.3, which is unusual, but not unheard of among developing countries. Nevertheless, the assumption of

[17] Specifically, by virtue of h being DRS in $\{L, K\}$, it follows that $\frac{\partial \ln h}{\partial \ln L} + \frac{\partial \ln Y}{\partial \ln K} < 1$ and $\frac{d\ln K}{d\ln L} < 1$.

Cobb-Douglas implies an elasticity of substitution of one between fixed and reproducible factors. If this elasticity is two[18] instead, Weil and Wilde report that halving population would cause a 26% increase in income per capita, or an elasticity of output to population of around -0.5.

While improving health means, in part, more population and more land dilution, this does not necessarily imply the return of Malthus in the 21st century. Some context is in order. What did the typical less-developed country look like 50 years ago versus today? The vast majority of its residents were in rural areas. Also, for most poor countries, this was the era of looking inwards as a growth strategy. Capital flows were sharply restricted, as was international trade. Migration to rich countries was also limited. All of these features accentuate the diminishing return to labor. But does this seem like a typical developing country today? Not really, because of two changes: urbanization and globalization.

Urbanization. The developing world is much less dependent on land than it was a half century ago. If the working-age population rises by x%, does the average plot of land get x% more farmers? Recent historical experience suggests not. As of a few years ago, the majority of the world's population lives in urban areas, continuing a long trend of migration to cities. Latin America, for example, went from 40% urban at mid-20th century to almost 80% today. Things today are less land-centered and more people-centered, with urbanization giving rise to agglomeration economies. In other words, we expect increasing rather than decreasing returns to scale. In the long run, land is less of a constraint, although congestion effects eventually kick in at much higher densities. (See Jena, Mulligan, Philipson, and Sun, 2007, for more analysis on the point.)

Globalization. The assumption of a closed economy seems increasingly unrealistic for the 21st century. How would openness change the results above? There are three mechanisms: capital flows, international trade, and migration. Allowing capital to flow across borders means that a country is not constrained by its own savings to finance investments. In sensitivity analysis, Ashraf, Lester, and Weil (2009) allow for capital to flow in from the outside world. This reduces capital dilution, and therefore output per capita falls less in the short run when population rises. Next, 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

[18] Weil and Wilde argue that the cross-country and time-series evidence is consistent with an elasticity of substitution of two. But it bears noting that estimates from the trend or cross section are not identified except under restrictive circumstances. Moreover, it is not clear whether historical experience can be a faithful guide to this issue today, for reasons that I discuss below.

US. (For most countries, however, this mechanism is probably of lesser importance that rural-to-urban migration.)

Openness to international trade should attenuate the effects of an expanding population on income per capita. Indeed, if the economy is diversified in the right way, we would expect factor-price equalization. This is seen in a simple model. Suppose that the economy can produce tradeable output in some sector (call it sector 1) that is not land dependent, but is instead CRS in labor and capital. Output in this sector is $Y_1 = g(K_1, L_1)$. Call sector 2 the land-dependent sector from above: $Y_2 = h(K_2, L_2, \bar{F})$. (Both g and h are standard, concave production functions.) The total stock of labor is $\bar{L}$, supplied inelastically, and $L_1 + L_2 = \bar{L}$. For an open economy, the prices of capital and the two goods are fixed in international markets at r, p_1, and p_2, respectively. Note that a competitive sector 1 has a perfectly elastic demand for labor at some wage $\tilde{w}$. Further, suppose the marginal product of labor in sector 2 is less than $\tilde{w}$ for $L_2 = \bar{L}$, so that both sectors produce something. If these two conditions hold, labor income per capita is no longer dependent on the stock of population.

It follows that policies that depress the non-land-dependent sector make the spectre of Malthus more likely. First, the above result holds only as long as sector 1 can expand. There may be some structural problem that checks the growth of sector 1 in the economy, which means that extra labor would go to sector 2 and income per capita would decline. Second, suppose that sector 1 is very unproductive for some reason. If the marginal product of labor in sector 2 at $L_2 = \bar{L}$ is greater than $\tilde{w}$, then sector 1 simply shuts down, and everyone produces in the land-dependent sector instead. If so, we are back in the Malthusian world modeled above, as if sector 1 did not exist. This analysis highlights that the effect of health on development interacts with other aspects of the economy.

But even if labor income per capita does not fall as $\bar{L} \uparrow$, there is one effect of population that is unavoidable: the dilution of land rents. The land-dependent sector above has decreasing returns to scale in reproducible factors, which creates rents for the owners of the fixed factors (or possibly for the government, which might expropriate these rents). As population grows, only the CRS sector expands, and so the land rents per capita decline. These rents can be especially large for extractive (rather than agricultural) activities, where the rents literally flow from the ground. As Weil and Wilde (2009, page 258) observe, "it is hard to imagine that Nigeria's oil production would be substantially different if the country had half (or double) its current population," and thus doubling the population would simply halve the oil revenue per capita.

Land Improvement. The threat of disease might deter land improvement. While we think of land as a 'fixed' factor, it is not in truly fixed (effective) supply in the sense that land can be improved. Malaria is often cited in this context. A leading example comes from the construction of the Panama Canal: the control of malaria on the isthmus was important for the successful completion of the canal by the United States, and conversely the failure of the French effort was partly due to a lack of scientific knowledge about malaria transmission at the time (McCullough, 1977). Relatedly, malaria control in Nepal spurred migration and land improvement in the lowlands of Nepal, although the resurgence of that disease in the 1980s undid some of gains from earlier control efforts (Seddon, 1995).

Compositional Changes. Beyond just population, health improvements will change the mix of ages and skills in the labor force. First, there is a direct effect of demographic composition on output: for example, by changing the fraction of the population that is of working age. But these demographic waves do not have permanent effects on the growth rate, absent some other structural change. Second, the changing composition of the population affects prices. Consider a few cases. First, suppose the shock to health brings more human capital to an economy where skilled labor is very scarce. In this case, the response of output would be more positive than the fixed-price 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 average income would rise less (or decline more) than expected.

Spillovers from Human Capital. External effects in production coming from human capital would complicate extrapolating from micro studies, which, by design, assume that the control group is not affected by the treatment variable. Measuring external effects, therefore, requires analyzing shocks at a higher level of aggregation (for example, the city or state) although there is no getting around the assumption that some area or group somewhere represents the counterfactual or control group. Additionally, there are considerable challenges to disentangling the causal external effect from correlations that arise from the sorting of factors across space. Moretti (2004) discusses the evidence of human-capital spillovers within US cities, mostly derived from cross-sectional and panel comparisons to estimate the effects. Acemoglu and Angrist (2000) use compulsory-schooling laws in the US as instruments for average education, and they find spillovers to income of 1%–2% per average year of schooling in one's state of residence, although they cannot rule out external effects of zero either. There are just a few studies directed at measuring external effects of human capital

in developing countries. Conley, Flyer, and Tsiang (2003) use geo-coded data to analyze the spatial aspects of human-capital spillovers in Thailand, although the properties of spatial equilibrium might not be informative about how the whole country would respond to an increase in human capital. Duflo (2004), on the other hand, analyzed the general-equilibrium impact of a school construction program in Indonesia. Duflo reports that the entrance of more educated cohorts appears to depress wages for incumbent, less educated cohorts. This suggests a negative spillover from human capital, at least at over some horizon.[19] In sum, the evidence is mixed, at best, for large, positive external effects of education in production[20] in developing economies.

Intergenerational effects. If having healthier parents raises the human capital of their children, the effects discussed above could be magnified. There is some evidence of intergenerational effects in the research mentioned above.[21] But how much does this add to the magnitude for income? Take an intergenerational elasticity of income to be 0.3. Should we simply multiply all of these effects by 1.3 to get the impact on the next generation? After all, if the nasty diseases are eradicated, the children's productivity rises by the benefit of growing up healthy along with 0.3 of the benefit of having healthier parents. Indeed, by this logic, the ith generation would receive 1.3^i higher income, which grows rather rapidly. For generations that are 25 years long, this adds about 1% to the per-capita growth rate of the economy. But it is unlikely that 0.3 is causal, at least in the sense that I apply it here. Furthermore, some of this transmission works through inputs that should be costed out, such as parents' time nurturing and children's time in school.

Morbidity versus Mortality. Not all improvements in health will increase population so automatically. For example, research on nutrition and parasitic disease find increases in productivity, but minor effects on mortality. This means that malnutrition and tropical parasites have their

[19] Because the incumbent cohorts were used as the control group by Duflo (2001), the c. 10% returns to schooling reported in that study are composed of both productivity and displacement effects, although we would want to exclude the latter in any calculation of the social return to schooling. Apart from social-versus-private considerations, the 10% number may be an overestmimate of the marginal return to school because any increase in school quality would have increased both the marginal and inframarginal returns to schooling. In terms of the model above, b_e would rise for various levels of e, not just at e^*, so simply dividing the sample averages of Δb by Δe would give a number greater than $\frac{db}{de}\Big|_{e^*}$.

[20] In contrast, there is ample evidence of spillovers in non-production activities. There are manifestly external effects of health: specifically the transmission of disease. But this effect is about the determination of average health rather than the effect of average health on income. There is also a large literature on peer effects in school, which I will not attempt to review here. Kremer and Holla (2009) discuss evidence of spillovers in school attendance, perhaps because of the consumption value of being with other kids. But this seems as likely to shift the choice of schooling up or down.

[21] Females given fresco instead of atole in the INCAP trial or who were in the womb for the Dutch winter famine both had smaller babies many years later (Behrman, 2009; Stein et al 1975).

greatest effect, not by adding more warm bodies to the labor force, but by preventing their victims from achieving their full productive potential.

5 MACRO ESTIMATES

5.1 Cross-Country Evidence

I could fill up my allotted space with references to papers that include some measure of health in a cross-country growth regression. (See Bloom, Canning, and Sevilla, 2004, for a partial review.) The predictive power of health variables is illustrated by Sala-i-Martin (1997) who reports which variables are most often statistically significant from literally millions of permutations. In this exercise, health variables (specifically life expectancy, malaria, and infant mortality) are among the most robust predictors of growth. (Note that this is not standard meta-analysis in that the sample for each test is the same.) Nevertheless, robustness in conditional correlation does not imply robustness in causation. Nor is it clear what failing this test would mean, insofar as health might affect output through other observable X variables. Conditioning on X would knock out health as a predictor, but this does not invalidate its role as a deeper cause.

It bears mentioning that these are probably, if anything, level rather than growth effects of health. First, although economic growth is on the left-hand side of growth/convergence regressions, the usual interpretation of these coefficients is that health affects levels.[22] Second, the micro estimates and human-capital theory above are also consistent with a level effect, although it is hard to imagine estimating a growth effect in a micro study, because the comparison groups are in the same economy as those treated.

Nevertheless, improving health could generate a temporary spurt of growth during the transition to a new level. This is especially important if it is early-life health that affects adult income. For instance, if the improvement in health benefits mostly infants, then none of the labor force is affected in the immediate aftermath, and this fraction increases only slightly for twenty years after the shock. Indeed, it would take more than half a century to realize the full effect on income if the mechanism were via exposure in infancy (Bleakley, 2006a).

A different set of studies relate economic output across countries to local malaria prevalence. Because the outcome variable is output per capita, this is a more direct method of assessing the

[22] A growth/convergence regression predicts economic growth as a function of the initial level of output per capita and other pre-determined variables. If the estimated system is stable, it implies a 'target' level of output per capita to which the country converges eventually. The right-hand-side variables determine what this target level is.

relation between health and output than the convergence regressions. Sachs (2003) reports that ln(GDP/capita) is lower by 1.3 as the fraction of the population potentially exposed to malaria goes from zero to one. Adjusted for differences in units, these estimates are about four times larger than the impact of persistent childhood exposure to malaria, even adjusting for general-equilibrium effects (Bleakley, 2006b and 2007b). Sachs (2003) also argues that health matters directly rather than simply through the historical determination of institutions. takes issue with the claim of Rodrik, Subramanian, and Trebbi (2004) that health variables only matter via the historical determination of institutions. Bhattacharyya (2009) finds that malaria is a robust determinant of output per capita in Africa data, and that controlling for malaria (or, indeed, climate) renders insignificant variables for institutional quality and the history of slave exports. (Indeed, he reports that simply controlling for frost is enough to knock out institutions or slave exports in Africa.) An advantage of these studies is that they use ecological conditions such as climate as an instrument for malaria, rather than the endogenous outcome of whether malaria has been eradicated. Nevertheless, although the ecology of malaria transmission is reasonably well modeled by local climate and geographic factors, it is an open question whether these factors affect economic activity only through malaria.

A few studies consider, at the country level, outcomes besides income. Lorentzen, McMillan, and Wacziarg (2008) relate adult mortality to increases in risky behavior. Bloom, Canning, and Graham (2003) show that longer life expectancy is associated with a higher saving rate. Bhattacharyya (2009) finds that saving is higher in African countries with less malaria. Alsan, Bloom, and Canning (2006) show that healthier countries receive more foreign direct investment.

Even if the causality could be resolved in the cross-country estimates discussed above, there remain difficult questions of when and how. Disease ecologies are persistent, and so a cross-sectional result does not clarify over what horizon health improvements would have their effects. The micro evidence above suggests that the very short-run effect is small, but even the potentially larger human-capital effects are an order of magnitude smaller than the macro estimates. Capital and land-improvement effects would also take effect slowly. (But the arguments above about the envelope theorem apply just as well to changing investments in capital and land.) The mechanism also matters for interpreting the result and designing health policies. Perhaps poor health is a root cause historically, but the general level of dysfunction that characterizes underdevelopment has co-evolved with the disease environment, and may impart a degree of hysteresis even if population health improves. If the effects of health work through human capital, then the educational system might need reinforcement; if the effects work through sectoral change, then laws might need to be

re-tooled; et cetera.

Furthermore, the 'how' and 'when' might interact in ways that render the result irrelevant for policy. Suppose we live in a world with some degree of path dependence. It is possible that health mattered at some critical juncture, sending the economy down one path or another. If this path got locked in for some reason, the effects of poor health would still be felt today, but not because health depresses income per se. A well-known variant of this idea is due to Acemoglu, Johnson, and Robinson (2001), who argue that the health environment influenced how European powers chose to colonize areas that came under their influence in the past 500 years, in particular regarding the imposition of extractive institutions. The choice of institutions then influenced the path of development, by their argument, which affects income above and beyond any direct effect of health. (Their claim is actually stronger than this: that most if not all of the effect of health on income works through the choice of colonial institutions.) Variants of this story that do not involve institutions are possible. For example, models of economic geography can imply path dependence as well. Health and human capital might have influenced the choice of industry at some crucial point in history (e.g. the Industrial Revolution), which could then have been locked in by agglomeration economies. (As noted by Smith (1905), this mechanism seems to have affected the distribution of population within tropical regions, which, for example, have numerous population centers and even national capital in sites on high plateaus that isolated but not malarial.) Put another way, even if it is completely correct to say that disease caused the current income to be lower, it cannot be inferred that treating disease will raise income by the same amount. To learn about the latter issue, we need to analyze the response of income to changes in health.

5.2 Response to Large Health Shocks

An alternative approach to this question is to analyze, at the aggregate level, the impact of large shocks to health that come from epidemics or eradication efforts. The advantage of analyzing such shocks is that the causality most likely runs from health to income. The downside, however, is that the results from a given shock may not generalize to other settings.

First, I discuss some regional studies that trace the impact of health shocks on output. A classic result is by Schultz (1964) who studied the 1918 influenza pandemic and farm output in India. The flu killed substantial numbers, but left capital and land intact, and had little scarring effect on the adults who survived. Schultz therefore interprets the flu mortality as a decline in

population only.[23] Comparing a few years before and after the epidemic (i.e., the relatively 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. Relatedly, I estimate the aggregate (i.e., state level) responses to the above-mentioned campaigns against hookworm and malaria in the United States (Bleakley, 2007b). Areas that stood to benefit from eradication saw slow increases in output that tracked the entrance into the workforce of cohorts exposed as children to the eradication campaigns. The estimated magnitudes are, if anything, larger than what I obtain in the earlier cohort-level estimates. This suggests little crowding of fixed factors.

These results highlight once again how the response of income to health depends on the manner in which health improves. The Schultz study is about mortality reducing the population, which increases the ratio of labor to land. The eradication of tropical parasites in the US South is, on the other hand, mostly about childhood morbidity. We would therefore expect it to work through human capital rather than fixed-factor dilution. This contrast is evidence that we need to unpackage health if we want to understand its impact on income.

Next, I turn to a recent study by Acemoglu and Johnson (2007, henceforth "AJ"), who consider the diffusion of technology *circa* 1950 that brought down mortality rates in poorer countries. AJ report that countries receiving the most favorable health shock experienced larger increases in life expectancy. With the increase in life expectancy came an increase in population, but a smaller and statistically insignificant rise in GDP. Thus the effect on output per capita was negative. They argue that other factors did not adjust when population grew.[24]

The magnitude of the AJ result is larger than what should be expected based on the work discussed above. Qualitatively, their result is similar to the one reported by Schultz, perhaps as expected because both studies are tilted towards mortality not morbidity. If anything, we would expect the Schultz estimate to exhibit more diminishing returns to labor than AJ's, because AJ look at longer horizons (decades instead of years), over which time 'sticky' factors should have adjusted. However, the magnitude of the AJ result is considerably larger: the ratio of the population to output elasticities (both w.r.t. life expectancy) is greater than 5. If we compare this to the bounds in Section 4.3, this implies a fixed-factor share of 0.8. This would be larger than Weil and Wilde's (2009) estimates of the land shares in income for any country today. The implied fixed-factor share would be larger still if the human-capital benefits of improved health are at all positive. Relatedly,

[23] This is not inconsistent with Almond's finding for the US in that the direct effects on productivity of *in utero* exposure to the pandemic would not be felt until those cohorts enter the labor force many years later.

[24] See Section 4.3 above for discussion of why this Malthusian interpretation of the episode may not be the best guide for health policy looking forward.

using a more general model, Ashraf, Lester, and Weil (2009, henceforth "ALW") report that the AJ result for GDP per capita is considerably worse than ALW's 'worst case' simulation.

As both AJ and ALW state, this muted response of output to population represents something of a puzzle. On the one hand, it may reflect compositional changes in the population. The reductions in infant mortality likely benefited the poor more than the rich, thus shifting the compositon of the population towards the children of the poor.[25] On the other hand, ALW suggest that it may reflect some yet-undiscovered specification problem.

In a recent comment, Bloom, Canning, and Fink (2009; henceforth "BCF") criticize AJ's first-stage regression by arguing it is sensitive to assumptions about mean reversion. BCF note high correlation between 1940 life expectancy and the change in life expectancy from 1940 to 2000. Consider the following bivariate regression using AJ's data:

$$\Delta e_{i,2000-1940} = \underset{(.022)}{-.753} \quad e_{i,1940} + \epsilon_{i,2000-1940} \tag{4}$$

where e is the natural log of life expectancy, and the standard error is in parentheses. This regression has an R^2 in excess of 0.96, which is a remarkable thing in and of itself. Perhaps not surprisingly, when BCF condition on initial life expectancy, this eliminates the AJ result by rendering the instrument insignificant in the first stage (which has $\Delta e_{2000-1940}$ on the left-hand side). BCF argue that controlling for mean reversion is needed.

Including the initial life expectancy as a control might seem innocuous at first, but there are at least two problems with doing so. First, note that a typical interpretation of equation 4 is convergence: countries with lower initial life expectancy tended to experience faster growth in life expectancy subsequently. But the convergence in life expectancy during this period might be precisely the variance that we would want to analyze. A large intervention that benefits unhealthy areas in this period would induce this relationship. (Indeed, some of the studies discussed above used pre-campaign health conditions $\times$ the date of the intervention to identify the effect of health.) Consistent with this idea, AJ report that 1940 life expectancy strongly predicts changes in life expectancy in the decades that follow, while 1970 life expectancy is a much weaker predictor of subsequent changes in life expectancy. Second, although some amount of convergence no doubt took place, we most likely cannot recover an unbiased estimate on e_{1940}, and thus conditioning on it directly would bias the other parameter estimates. In the 1940s, many countries did not have

[25] The micro literature suggests that scarring effects are larger than selection effects, but the filter of publication bias might mean that we never see studies in which selection dominates scarring.

adequate vital-registration systems. (Even the United States did not have a nationwide system of death registration until the 1930s.) Suppose that the true model for equation 4 is $e_1^* - e_0^* = \beta e_0^* + \epsilon_1^*$ but that we observe e_t, which is the truth plus error: $e_t = e_t^* + \eta_t$. If the measurement error is 'classical', meaning that $cov(e_0, \eta_t) = 0$ for $t = 0, 1$, then plim $\hat{\beta} = \theta\beta + (1-\theta)(-1)$, for $\theta \equiv \frac{\text{var}e_0^*}{\text{var}e_0^* + \text{var}\eta_0}$. As the variance of η_0 increases, this estimate is biased towards -1.

Measurement error could also affect the AJ result. For long differences, the AJ instrument can be written as

$$z_0^* = \ln \Sigma_i \left(\frac{\tilde{M}_{j0}^*}{\tilde{P}_0^*} \right) = \tilde{m}_0^* - \tilde{p}_0^*,$$

where $\tilde{M}_{j0}^*$ is baseline mortality for disease j multiplied by the predicted decline based on global advances in health technology, $\tilde{m}_0^*$ is a aggregate over all the diseases j, and P and p are population, which is used to normalize the disease rates. The tildes on mortality and population variables reflect the fact that AJ used cities in some cases to proxy for the country-level data. Capital letters denote levels and lower-case letters are natural logs. The reduced-form and first-stage equations are, respectively, $\Delta p_1^* = \gamma z_0^* + \varepsilon_1^*$ and $\Delta e_1^* = \alpha z_0^* + \omega_1^*$, where AJ assume the error terms are uncorrelated with z_0^*. Now suppose that we also observe population and baseline mortality with error, which seems more than likely.[26] Let measured population and baseline mortality variables be $p_t = p_t^* + \xi_t$. and $m_0 = m_0^* + \upsilon_0$. We expect, then, the instrumental variables (IV) estimate for the effect of life expectancy on population to converge to

$$\text{plim } (\hat{\gamma}/\hat{\alpha}) = \frac{\text{cov}\left(\Delta p_1, z_0\right)}{\text{cov}\left(\Delta e_1, z_0\right)} = \frac{\text{cov}\left(\gamma z_0^* + \varepsilon_1^* + \xi_1 - \xi_0 \ , \ z_0^* + \tilde{\upsilon}_0 - \tilde{\xi}_0\right)}{\text{cov}\left(\alpha z_0^* + \omega_1^* + \eta_1 - \eta_0 \ , \ z_0^* + \tilde{\upsilon}_0 - \tilde{\xi}_0\right)},$$

where, again, the tildes reflect the possible use of sub-national data to proxy for a country's cause-specific mortality rate. This probability limit does not, in general, go to γ/α.

Consider the sort of measurement error that arises in a few salient cases. First, census under-enumeration was probably greater in poor countries, especially in the early part of the sample. The ξ_0 would be correlated with GDP/capita, life expectancy, and (inversely with) mortality, and $\xi_1 - \xi_0$ would therefore be higher in less-developed countries. This would bias $\hat{\gamma}/\hat{\alpha}$ upwards. Second, estimates of total mortality, life expectancy, and population growth were not generated independently, especially in underdeveloped countries. (Indeed, in some cases the life expectancy

[26] Measurement error in population is probable because (i) census coverage in this period was incomplete for many countries and (ii) some of the countries in AJ's sample data did not actually have censuses in the baseline year (Durand, 1950). Errors in enumeration of 5%–15% were not uncommon in poor countries circa 1940. The case of the Guatemala census of this period is notorious: it was believed to be off by one third.

numbers were imputed from infant mortality, changes in population, and model life tables.) By construction, measurement errors would be correlated across variables. Specifically, $\text{cov}(\eta_0, \upsilon_0,) < 0$ and $\text{cov}(\xi_1 - \xi_0, \upsilon_0) < 0$, which would could bias $\hat{\gamma}/\hat{\alpha}$ up or down. Third, the use of sub-national proxies of mortality for some countries mean that mortality rates are mismeasured. This particular measurement error might just be classical, which would not bias the IV estimate. Moreover, the use of an alternate proxy of mortality is useful in that it attenuates the division bias that would have resulted from including the same (mis)measure of baseline population on both sides of the reduced-form equation for population. (That is, any upward bias in the IV estimate is attenuated because presumably the $\text{cov}(\xi_0, \tilde{\xi}_0)$ is less than $\text{var}(\xi_0)$.)

How much would these measurement problems matter? It is hard to know without further understanding of the error properties of these data. The AJ result is some of the best evidence we have using cross-country comparisons, but further research is needed into whether reliable inferences can be made with such data.

6 CONCLUSIONS

This review is organized around the following question: how much does poor health hold back development? I focus on human capital and income. The micro evidence finds that childhood health is an input to producing other forms of human capital and that adult health depresses productivity as well. I use a standard model of human capital to tie these various results together. Then, applying the Envelope Theorem to the model, I argue for a re-orientation of the micro literature; in short, we should be looking at outputs like income rather than inputs like time in school.

Next, I review the aggregate implications of micro estimates, which are complicated by health's effect on population size and other general-equilibrium issues. I argue that further research is needed in understanding how health improvements interact with other aspects of the economy, as well as how fertility responds to health shocks. I also review the macro evidence on this topic, which has taken the form of either cross-country comparisons or measuring the responses to large health shocks.

I conclude by reflecting on why we study interrelationships between health and development in the first place. Is the question whether health is the single factor that explains everything, leading to a 'magic bullet' intervention? This sets a pretty high standard, one that health (along with everything else) fails. Are we instead looking for do-able interventions that pass the cost/benefit

test? If so, there seem to be such examples in disease control.

References

Acemoglu, Daron and Joshua Angrist (2000). "How Large Are Human-Capital Externalities? Evidence from Compulsory Schooling Laws." in B. S. Bernanke and K. Rogoff, eds., *NBER Macroeconomics Annual*, NBER MIT Press Cambridge, Massachusetts November pp. 9–58.

——— and Simon Johnson (2007). "Disease and Development: The Effect of Life Expectancy on Economic Growth." *Journal of Political Economy.* December **115** (6), 925–985.

——— , ——— , and James A. Robinson (2001). "The Colonial Origins of Comparative Development: An Empirical Investigation." *American Economic Review.* December **91** (5), 1369–401.

Almond, Douglas (2006). "Is the 1918 Influenza Pandemic Over? Long-term Effects of In Utero Influenza Exposure in the Post-1940 U.S. Population." *Journal of Political Economy.* August **114** (4), 672–712.

——— and Bhashkar Mazumder (2008). "Health Capital and the Prenatal Environment: The Effect of Maternal Fasting During Pregnancy." Working Paper 14428, National Bureau of Economic Research, Cambridge, Mass. October.

Alsan, Marcella, David E. Bloom, and David Canning (2006). "The effect of population health on foreign direct investment inflows to low- and middle-income countries." *World Development.* April **34** (4), 613–630.

Ashraf, Quamrul H., Ashley Lester, and David N. Weil (2009). "When Does Improving Health Raise GDP?" *NBER Macroeconomics Annual.*

Barlow, Robin (1967). "The Economic Effects of Malaria Eradication." *American Economic Review.* May **57** (2), 130–148.

Barreca, Alan (2007). "The Long-Term Economic Impact of *In Utero* and Postnatal Exposure to Malaria." Technical Report, UC-Davis.

Behrman, Jere R. (2009). "Early Life Nutrition and Subsequent Education, Health, Wage, and Intergenerational Effects." in M. Spence and M. Lewis, eds., *Health and Growth*, Washington, DC: Commission on Growth and Development and the World Bank, chapter 6, pp. 167–183.

——— and Mark R. Rosenzweig (2004). "Returns to Birthweight." *Review of Economics and Statistics.* **86** (2), 586–601.

Bhattacharyya, Sambit (2009). "Root Causes of African Underdevelopment." *Journal of African Economies.* **10.**

Bleakley, Hoyt (2006a). "Disease and Development: Comments on Acemoglu and Johnson (2006)." Delivered at NBER Summer Institute, EFG Meeting July. http://home.uchicago.edu/∼bleakley.

——— (2006b). "Malaria in the Americas: A Retrospective Analysis of Childhood Exposure." Documento CEDE 2006-35, Universidad de los Andes, Bogotá, Colombia. September.

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

——— (2007b). "Spillovers and Aggregate Effects of Health Capital: Evidence from Campaigns Against Parasitic Disease in the Americas." Unpublished manuscript, University of Chicago.

——— and Fabian Lange (2009). "Chronic Disease Burden and the Interaction of Education, Fertility and Growth." *Review of Economics and Statistics.* February **91** (1), 52–65.

Bloom, David E., David Canning, and Bryan Graham (2003). "Longevity and Life-cycle Savings." *Scandinavian Journal of Economics.* 09 **105** (3), 319–338.

——— , ——— , and Gunther Fink (2009). "Disease and Development Revisited." NBER Working Papers 15137, National Bureau of Economic Research July.

Bobonis, Gustavo J., Edward Miguel, and Charu Puri Sharma (2006). "Anemia and School Participation." *Journal of Human Resources.* **41** (4), 692–721.

Brinkley, Garland L. (1997). "The Decline in Southern Agricultural Output, 1860-1880." *Journal of Economic History.* March **57** (1), 116–38.

Bundy, Donald A. P., Michael Kremer, Hoyt Bleakley, Matthew C. H. Jukes, and Edward Miguel (2009). "Deworming and Development: Asking the Right Questions, Asking the Questions Right." *PLOS/NTD.* January **3** (1).

Conley, Timothy G., Fredrick Flyer, and Grace R. Tsiang (2003). "Spillovers from Local Market Human Capital and the Spatial Distribution of Productivity in Malaysia." *Advances in Economic Analysis and Policy.* **3** (1).

Cunha, Flavio and James J. Heckman (2008). "Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation." *Journal of Human Resources.* **43** (4).

Cutler, David, Winnie Fung, Michael Kremer, Monica Singhal, and Tom Vogl (2007). "Mosquitoes: The Long-term Effects of Malaria Eradication in India." Working Paper 13539, National Bureau of Economic Research October. (Revised on August 17, 2009.).

Delange, François M. and Jr. Keith P. West, eds (2003). *Micronutrient deficiencies in the first months of life*, Vevey, Switzerland: Nestlé Nutrition.

Duflo, Esther (2001). "Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment." *American Economic Review.* September **91** (4), 795–813.

——— (2004). "The Medium Run Effects of Educational Expansion: Evidence from a Large School Construction Program in Indonesia." *Journal of Development Economics.* June **74** (1), 163–197.

Durand, John D. (1950). "Adequacy of Existing Census Statistics for Basic Demographic Research." *Population Studies.* September **4** (2), 179–199.

Fortson, Jane (2009). "Mortality Risk and Human Capital Investment: The Impact of HIV/AIDS in Sub-Saharan Africa." August. Manuscript, Mathematica Policy Research.

Glewwe, Paul and Elizabeth M. King (2001). "The Impact of Early Childhood Nutritional Status on Cognitive Development: Does the Timing of Malnutrition Matter?" *The World Bank Economic Review.* September **15** (1), 81–113.

Hanson, Gordon H. and Craig McIntosh (2007). "The Great Mexican Emigration." Working Paper 13675, National Bureau of Economic Research December.

Hazan, Moshe (2006). "Longevity and Lifetime Labour Input: Data and Implications." Discussion Papers 5963, CEPR November.

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

Jayachandran, Seema and Adriana Lleras Muney (2009). "Life Expectancy and Human Capital Investments: Evidence from Maternal Mortality Declines." *Quarterly Journal of Economics.* February **124** (1), 349–397.

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

Joshi, Shareen and T. Paul Schultz (2007). "Family Planning as an Investment in Development: Evaluation of a Program's Consequences in Matlab, Bangladesh." Working Paper, Economic Growth Center, Yale University February.

Kalemli-Ozcan, Sebnem, Harl E. Ryder, and David N. Weil (2000). "Mortality decline, human capital investment, and economic growth." *Journal of Development Economics.* **62** (1), 1–23.

Kremer, Michael and Alaka Holla (2009). "Improving Education in the Developing World: What Have We Learned from Randomized Evaluation." *Annual Reviews of Economics.* **I** (19).

Lorentzen, Peter, John McMillan, and Romain Wacziarg (2008). "Death and development." *Journal of Economic Growth.* June **13** (2), 81–124.

Lucas, Adrienne M. (2009). "Malaria Eradication and Educational Attainment: Evidence from Paraguay and Sri Lanka." February. Unpublished manuscript.

Maccini, Sharon and Dean Yang (2009). "Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall." *American Economic Review.* June **99** (3), 1006–26.

McCullough, David (1977). *The Path Between the Seas: The Creation of the Panama Canal, 1870-1914*, New York: Simon and Schuster.

Meng, Xin and Nancy Qian (2006). "The Long Run Health and Economic Consequences of Famine on Survivors: Evidence from China's Great Famine." Discussion Papers 5989, CEPR 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.

Miller, N. Grant (2005). "Contraception as Development? New Evidence from Family Planning in Colombia." Working Paper 11704, National Bureau of Economic Research, Cambridge, Mass.

Moretti, Enrico (2004). "Human Capital Externalities in Cities." in J. V. Henderson and c. T. Jacques-Fran eds., *Handbook of Regional and Urban Economics*, Vol. 4, Amsterdam: North-Holland, chapter 51, pp. 2243–2291.

Philipson, Tomas J. (2000). "External Treatment Effects and Program Implementation Bias." Working Paper T0250, National Bureau of Economic Research, Cambridge, Mass. January.

Preston, Samuel H. (1980). "Causes and consequences of mortality declines in less developed countries during the twentieth century." in R. Easterlin, ed., *Population and Economic Change in Developing Countries*, Chicago, Illinois: National Bureau of Economic Research, chapter 5, pp. 289–359.

Ribero, Rocío and Jairo Nuñez (2000). "Adult Morbidity, Height, and Earnings in Colombia." in W. D. Savedoff and T. P. Schultz, eds., *Wealth from Health: Linking Social Investments to Earnings in Latin America*, Washington, D.C.: Inter-American Development Bank.

Rodrik, Dani, Arvind Subramanian, and Francesco Trebbi (2004). "Institutions Rule: The Primacy of Institutions over Geography and Integration in Economic Development." *Journal of Economic Growth.* June **9** (2), 131–165.

Rosenzweig, Mark R. and Kenneth I. Wolpin (1980). "Testing the Quantity-Quality Fertility Model: The Use of Twins as a Natural Experiment." *Econometrica.* **48** (1), 227–240.

Sachs, Jeffrey D. (2003). "Institutions Don't Rule: Direct Effects of Geography on Per Capita Income." Working Paper 9490, National Bureau of Economic Research, Cambridge, Mass. February.

Sala-i-Martin, Xavier (1997). "I Just Ran Two Million Regressions." *American Economic Review.* May **87** (2), 178–83.

Savedoff, William D. and T. Paul Schultz, eds (2000). *Wealth from Health: Linking Social Investments to Earnings in Latin America*, Washington, D.C.: Inter-American Development Bank.

Schapiro, Louis (1919). "The Physical and Economic Benefits of Treatment for Hookworm Disease." *Journal of the American Medical Association.* **73** (20), 1507–1509.

Schultz, T Paul (1985). "Changing World Prices, Women's Wages, and the Fertility Transition: Sweden, 1860-1910." *Journal of Political Economy.* December **93** (6), 1126–54.

Schultz, Theodore W. (1964). *Transforming Traditional Agriculture* number 3. In 'Studies in Comparative Economics.', New Haven, CT: Yale University Press.

Seddon, David (1995). "Migration: Nepal and India." in R. Cohen, ed., *The Cambridge survey of world migration*, Cambridge, UK: Cambridge University Press, pp. 367–372.

Shoda, Y., W. Mischel, and P. K. Peake (1990). "Predicting adolescent cognitive and social competence from preschool delay of gratification: Identifying diagnostic conditions." *Developmental Psychology.* **26**, 978–986.

Smith, J. Russell (1905). "The Economic Importance of the Plateaux in Tropic America." *Bulletin of the American Geographical Society.* **37** (8), 461–468.

Stein, Zena, Mervyn Susser, Gerhart Saenger, and Francis Marolla (1975). *Famine and Human Development: The Dutch Hunger Winter of 1944-1945*, London: Oxford University Press.

Thomas, Duncan, Elizabeth Frankenberg, Jed Friedman, Jean-Pierre Habicht, Mohamme Hakimi, Jaswadi, Nathan Jones, Gretel Pelto, Bondan Sikoki, Teresa Seeman, James P. Smith, Cecep Sumantri, Wayan Suriastini, and Siswanto Wilpo (2003). "Iron deficiency and the well-being of older adults: Early results from a randomized nutrition intervention." April. Unpublished manuscript.

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

——— and Joshua Wilde (2009). "How Relevant Is Malthus for Economic Development Today?" *American Economic Review.* May **99** (2), 255–60.

Weisbrod, Burton A., Ralph L. Andreano, Robert E. Baldwin, Erwin H. Epstein, and Allen C. Kelley (1973). *Disease and Economic Development: The Impact of Parasitic Diseases in Saint Lucia*, Madison, Wisconsin: The University of Wisconsin Press.

Young, Alwyn (2005). "The Gift of the Dying: The Tragedy of AIDS and the Welfare of Future African Generations." *Quarterly Journal of Economics.*