When Does Improving Health Raise GDP?

Quamrul H. Ashraf, Ashley Lester, and David N. Weil*
Brown University
Department of Economics

Very preliminary and incomplete. Please do not cite without permission.

April 20, 2007

Abstract

We assess quantitatively the effect of exogenous health improvements on output, through demographic channels and changes in worker productivity. We consider both changes in general health, proxied by changes in life expectancy, and changes in the prevalence of particular diseases, such as malaria and tuberculosis. In general, we find that the effects of health improvements on income are substantially lower than those that are often quoted by policy-makers, and may not emerge at all for half a century or more after the initial improvement in health.

*We would like to thank seminar participants at the Harvard Program on Global Demography and Aging and the Brown Macro lunch for their helpful suggestions. We are grateful to the William and Flora Hewlett Foundation for financial support. Comments are welcome to the authors at quamrul_ashraf@brown.edu, ashley_lester@brown.edu, or david_weil@brown.edu.

1 Introduction

What are the long-term economic outcomes of interventions that affect population health? Can disease eradication raise per capita income in the developing world, or will it lead merely to a population explosion and the attendant pressures on resources and income? If improving health can raise per capita income, how long might the beneficial effects take to manifest themselves? This paper investigates the effects of changes in health on macro-economic outcomes, and highlights the role of demographics in transmitting these changes.

A strong consensus exists among economists that an individual's health is an important determinant of his or her economic performance. Various measures of poor health, including malnutrition, anemia and exposure to disease *in utero* and during childhood, have all been shown to have a negative effect on a person's wages or productivity. Moreover, at the macroeconomic level, there is a strong positive correlation between income per capita and life-expectancy or other measures of health. Based on this evidence, it is often argued that a major reason why some countries are less developed is the poor health of their populations, and therefore that major improvements in health would cause significant economic gains. For example, the report of the **?**, chaired by Jeffrey Sachs, finds evidence that health is one of the most important determinants of a country's economic success. Similarly, the Copenhagen Consensus 2004, a gathering of distinguished economists including four Nobel laureates, ranked interventions against HIV/AIDS and malaria the first and fourth-highest priorities respectively for world development. On humanitarian grounds alone, this conclusion is difficult to fault. But the supporting documents do not confine their arguments to humanitarian principles, arguing strongly that "malaria and HIV/AIDS are major causes of disease burden and economic losses" (**?**). These arguments have found a receptive audience among policy-makers. The Abuja Declaration of 2005, signed by fifty-three African heads of state, notes that "malaria has slowed economic growth in African countries by 1.3% per year as a result of which GDP for African countries is now 37% lower than it would have been in the absence of malaria."

However, the argument for drawing a macroeconomic conclusion from either the microeconomic evidence or the cross-sectional correlation is problematic. Most importantly, microeconomic studies are unable to control for the general equilibrium effect of changes in population health. For example, an increase in life expectancy may lead to a larger population, in turn reducing available per capita resources and possibly undoing the economic benefits of better health. On the other hand, macroeconomic cross-country regressions that could potentially capture these effects typically suffer from omitted variables bias and reverse causation problems. In an important recent study, **?** find that, when the problems of health's endogeneity and omitted variables are corrected, health improvements in the period after World War II actually had a negative effect on income per capita.

This paper revisits the question of the extent to which health determines national income and asks by how much national income can be raised by interventions that improve health. We use a simulation model to generate long-run macroeconomic predictions building up from microeconomic

estimates. Our model requires the specification of three sets of parameters: those involving the demographic response to health and mortality changes; those involving the effect of health on labor productivity and other aspects of human capital; and those involving the aggregate production function.

We apply the model to two distinct types of exogenous changes in health. The first is an increase in life expectancy, treating life expectancy as a summary measure of the general state of health in a nation. In particular, we consider the effect of exogenously raising life expectancy at birth from 40 to 60 years. This approximately corresponds with the most dramatic improvement in health observed during the international epidemiological transition studied by **?**. The key finding from these simulations is that even large increases in life-expectancy, which could raise per capita income in the long-run by around 15 per cent, may reduce income by up to 10 per cent for fifty years or more after the shock.

The second type of change in health we consider is the exogenous eradication of a particular disease or set of diseases. At present, our results focus on two infectious diseases that are particularly prevalent in the developing world; malaria and tuberculosis. These simulations have two key results. The first is that, for either of the diseases considered, even complete eradication has a relatively small impact on income per-capita in either the short or the long-run, not exceeding a few per cent of GDP. The second is that these relatively small effects do, however, vary quite widely by disease. For example, in the short-run, eradicating tuberculosis raises income per-capita whereas eradicating malaria lowers it. In the long-run, eradicating tuberculosis has a small positive effect on GDP per capita, while eradicating malaria is approximately neutral in its effects on GDP per capita. The different effects on income of eradicating these diseases arise largely because tuberculosis strikes relatively more prime-age workers, while malaria affects relatively more young children.

The simulation-based methodology allows us to take into account both general equilibrium effects and the dynamic effect of health through channels including the evolution of the size and age-structure of the population, capital accumulation and resource depletion. The analysis is well-adapted to considering the dynamic path of the economy over the course of this evolution, rather than merely comparing steady-states, and to providing a quantitative characterization of this evolution in the face of particular interventions, such as control of malaria or tuberculosis. The simulation approach also permits analysis of the strength of the various mechanisms at work. For instance, it is straightforward to examine the sensitivity of the results to different estimates of the effect of disease on effective labor supply or the speed of the demographic transition. Although the results appear quite robust to alternative specifications of the direct effect of health on human capital, they are sensitive to changes in assumptions about the timing of fertility reduction.

Our results are likely to be of interest to policy-makers in developing countries, as illustrated by the Abuja declaration. Although income per capita is an inadequate measure of welfare, particularly when considering changes in health or fertility, it is nevertheless of widespread interest.

Moreover, a quantitative analysis of the effects of health interventions on the economy may be useful in determining the cost and priority of a range of policies, including the economic benefits of competing possible health interventions and the relative importance of provision of family-planning services in association with interventions aimed at raising health.

Section Two of the paper presents the structure of the model and the parametric choices underlying our results. Section Three presents the results of our alternative health experiments, and considers the sensitivity of the results to alternative methods of measuring the effect of health on labor supply. Section Four presents some sensitivity analysis to investigate how robust the results are to changes in assumptions about experience, resource depletion and fertility. Section Five concludes and raises some suggestions for future research.

2 The Model and its Parameterization

As stressed above, there are a number of channels through which health affects income, and the dynamics of these effects can be stretched out over several decades. Thus, analyzing the effect of a health intervention entails comparing the complete paths of income and other endogenous variables in the scenario in which the intervention takes place to an alternative in which it does not. Similarly, alternative paramaterizations of the different components of the model will yield different dynamic paths of all the endogenous variables.

We consider two different sorts of health interventions. First, we consider a "general" health improvement. Specifically, we consider a shift in life-expectancy at birth from $e_0 = 40$ to $e_0 = 60$ using a model life table. Second, we consider the eradication of two specific diseases: malaria and tuberculosis.

The model features both demographic and economic elements. The demographic elements comprise estimates of mortality and fertility by age. The economic elements include the specification of the aggregate production function and the specification of the response of variables such as human capital to changes in health. We consider each element in turn.

2.1 Demographic Structure

The demographic part of the model takes age-specific mortality and fertility schedules as inputs to create pre- and post-intervention population projection matrices. We first use the pre-intervention matrix to generate a stable population, and then introduce a permanent shock by applying the post-intervention matrix (incorporating a different mortality schedule) to project the population forward. For example, the shock could involve a shift in life-expectancy at birth from $e_0 = 40$ to $e_0 = 60$ in a model life table or, if we are considering the eradication of a specific disease, we can use a table of age-specific deaths from that disease to create a cause-deleted mortality schedule. Formally, a population composed of n age-groups is represented by an n-dimensional vector, N_t,

for every period t, and evolves according to:

$$N_{t+1} = \begin{cases} P^b \cdot N_t & \text{if } t < T \\ \\ P^a \cdot N_t & \text{otherwise,} \end{cases} \tag{1}$$

where P^b and P^a are the $n \times n$ projection matrices before and after the shock, $N_0 > 0$ is given, and the shock period, T, is determined to occur after the pre-shock population has attained a stable age structure and rate of growth.[1] In practice, population is divided into approximately 5-year age groups. We assume that children enter the labor force at 15 and workers leave the labor force at 65.

The critical demographic parameters involved are those that govern the evolution of the population after the shock. In particular, we need to make a long-run assumption about how fertility rates change in response to the change in life-expectancy, and a short-run assumption about the speed with which fertility rates adjust to their new long-run level.

For the long-run, we assume that the growth rate of the population eventually returns to its pre-shock level. In the short-run, fertility responds to mortality only with a lag, and so we update P^a accordingly over a specified number of periods to incorporate the demographic transition. The demographic part of the model thus generates a dynamic path for both population size and age structure that is broadly consistent with the stylized characterization of the demographic transition.

Estimates of the appropriate lag to apply to the transition period are difficult to come by, so we currently assume as a base case that the growth rate of the population returns to its pre-shock level over a period of 50 years after the shock. This is a little slower than some accounts of the East Asian demographic transition, which suggest a period of about 25-30 years, but seems quite consistent with evidence from Europe and India and perhaps even a little optimistic in Africa (see **?** for a brief summary of the evidence). We can also test the sensitivity of our results to different assumptions regarding the length of the transition period.

2.2 Production and Physical Capital Accumulation

Aggregate production is modelled using a standard Cobb-Douglas production function. The factor inputs are land, capital and human capital, so that aggregate output in period t, Y_t, is:

$$Y_t = A_t K_t^{\alpha} H_t^{\beta} X^{1-\alpha-\beta} \tag{2}$$

where $\alpha \in (0,1)$, $\beta \in (0,1)$, $\alpha + \beta \leq 1$. X is a fixed arbitrary stock of land and A_t is a Solow residual.

[1] Specifically, a population projection matrix is composed of age-specific net maternity rates along the first row and age-specific survivorship rates along the sub-diagonal. The stable population growth rate implied by a projection matrix is given by its largest, real eigenvalue, and the stable age-structure by the corresponding eigenvector.

We assume fairly standard values for factor shares: we set $\alpha = 0.3$ and $\beta = 0.6$, meaning that the implied fixed share of land is 10%. In future work we will use data from the World Bank on shares of natural capital in national income at the country level to allow us to estimate a somewhat more flexible functional form for the aggregate production function.[2]

The Solow residual grows at an exogenous rate that does not respond to any of the changes in the model. For convenience, the growth rate is set to equal the steady-state rate of population growth, so that income per-capita is constant in the steady-state. Because all of our results entail a comparison of income in the case of a health intervention to the case where no intervention takes place, the underlying rate of technological change is of very little importance.

We handle capital accumulation extremely simply, by making the Solovian assumption that a fixed share of national income is saved in each period.[3] Accordingly, the stock of capital in period t, K_t, evolves over time according to:

$$K_{t+1} = sY_t + (1 - \delta)K_t, \tag{3}$$

where $s \in (0, 1)$ and $\delta \in (0, 1)$ are the fixed saving and depreciation rates, respectively, and $K_0 > 0$ is given. We assume that the savings rate is 10 per cent, which is close to the average for countries in sub-Saharan Africa, and assign a standard value to the depreciation rate such that $\delta = 0.08$.

2.3 Human Capital

We model an individual's human capital as a function of his or her schooling, experience, and health. We assume that human capital inputs of individuals with different characteristics are perfectly substitutable. Thus the stock of human capital in period t, H_t, is the product of human capital per worker (from health, schooling and experience) across all individuals in the working age groups in period t. In particular,

$$H_t = \sum_{15 \leq i \leq 65} \left(h^h_{i,t} \times h^s_{i,t} \times h^e_{i,t} \right) N_{i,t}, \tag{4}$$

where $N_{i,t}$ is the number of individuals of age i in the population in period t.

Our treatment of schooling and experience is relatively standard. Human capital from schooling is constructed using estimates of years of schooling from the **?** dataset. Years of schooling

[2] Another approach would be to compare natural capital (capitalized value of subsoil resources, farm and pasture land, etc.) to produced physical capital, on a per capita basis.

[3] **?** makes the same assumption in his analysis of HIV/AIDS in South Africa. An alternative would be to build in a life-cycle model of saving, although there is considerable controversy about the applicability of such models to developing countries. See **?** and **?**.

are then aggregated into human capital from schooling using the piecewise log-linear specification:

$$h^s_{i,t} = \begin{cases} \exp(\theta_1 S) & \text{if} \quad S \leq 4 \\ \exp(4\theta_1 + \theta_2 (S-4)) & \text{if} \quad 4 < S \leq 8 \\ \exp(4\theta_1 + 4\theta_2 + \theta_3 (S-8)) & \text{if} \quad S > 8 \end{cases}$$

We use values of $\theta_1 = 0.134$, $\theta_2 = 0.101$, and $\theta_3 = 0.068$, based on **?**.

Human capital from on-the-job experience for a worker of age i in any period t, $h^e_{i,t}$, is computed as:

$$h^e_{i,t} = \exp\left[\phi(i-15) + \psi(i-15)^2\right] \tag{5}$$

where, based on **?** who provide an estimate of the average return to experience in a sample of 48 countries, we use a ϕ value of 0.0495 and a value of -0.0007 for ψ.

2.3.1 Human Capital from Health

As mentioned above, we consider two different types of health interventions: a general improvement in health (an increase in life expectancy at birth from 40 to 60), and the eradication of specific diseases. We further use two different methods for parameterizing the effects of a general health improvement.

The first method for modeling the effect of a general improvement in health uses the estimate of the effect of increased adult survival rates (ASR) on productivity generated in **?**, which in turn draws on a large number of well-identified microeconomic studies. These studies have generated considerable evidence that health interventions during adulthood can beneficially affect worker productivity. The ASR is defined as the probability that an individual will attain the age of 60, conditional on having attained the age of 15. The preliminary results of our simulations reported in the next section arise from employing a ρ value of 0.653, based on estimates in **?**.

To give a concrete example of the size of this effect, a change in life expectancy at birth from 40 to 60 corresponds, using the UN female model life table for the South Asia region, to a change in the ASR from 0.50 to 0.72. Applying the coefficient above implies an increase of 15% in health human capital per worker.

The second method of capturing the direct effect of health improvements on worker productivity uses the ratings of disease incidence and severity that are used to construct estimates of years lost due to disability (YLD) around the world by the World Health Organization (WHO). The WHO provides a general measure of YLDs and then also measures disease specific YLDs, both broken down by age group. A country's YLD for a given disease is constructed as:

$$YLD = I \times DW \times L$$

where I is the number of incident (newly-arising) cases in a period, DW is the disability weight attached to the disease, and L is the average duration of the disease until remission or death.

The crucial parameter here is the disability weight, which is intended to be a cardinal measure of the severity of different diseases or impairments, on a scale from 0, indicating perfect health, to 1, indicating death. Disability weights are constructed largely on the basis of interviews with households, who are asked to value the various health states in question. They are therefore not primarily intended as a measure of labor supply. Nevertheless, these estimates provide at least some basis for comparing the effect of different diseases, as well as a cross-check on the results using the ASR parameter discussed earlier.

Because YLD data play a significant role in the analysis below, it is worth exploring these data in more detail. Table 1 shows data from the WHO "AFRO E" region, (defined as Africa with high child and very high adult mortality). We look at men in the 30-44 age group, and show data on both total population and YLDs as well as per capita YLDs. Overall, men in this group average 13.5% of a YLD per capita per year, with one third of this burden coming from infectious and parasitic diseases. HIV/AIDS makes up half of the infectious disease burden, while the two diseases that we consider below make relatively small contributions. Tuberculosis accounts for .005 YLDs, or 3.5% of the disability burden, while malaria accounts for only .001 YLDs, or 1% of the total disability burden. The disability weight for tuberculosis is .271. For an episode of malaria, the disability weight is .191, while the disability weight for the neurologic sequelae of malaria is .471.

To assess the effect on worker productivity from a general health improvement (i.e. an increase in life expectancy at birth from 40 to 60) we need a mapping from life expectancy to YLDs. We construct this mapping by looking at cross sectional data from 14 WHO sub-regions on YLDs per capita and life expectancy at birth. Figure 1 shows the data for all age groups. In practice, we work with similar data at the age-group level (each group spans approximately 15 years for most of the working-age population). For each age group, we run a regression across the 14 sub-regions of YLDs per capita on life expectancy at birth. The coefficients from these regressions then tell us the change in age-group specific YLDs that would result from an increase in life expectancy from 40 to 60. To give a concrete example, the regression of YLDs per capita on life expectancy at birth for the 30-44 age group (with standard errors in parentheses) is

$$\begin{array}{rcc} YLD = & .251 & -.0023e_0, \quad r^2 = .88 \\ & (.017) & (.0002) \end{array}$$

Applying these regression coefficients, a change in life expectancy at birth from 40 to 60 would lower YLDs from .159 to .113, implying a rise in labor input of 5.5 per cent. Similarly, the implied increases in labor input per worker in the 15-29 and 45-59 age groups are 4.5 per cent and 5.0 per cent, respectively. These effects are roughly one-third the size of the effects we estimate using the data on ASR.

Conceptually, both the ASR and YLD estimates are derived from thinking about a comparison of workers who have spent their entire lives in a low- or high-life expectancy environment.

However, in response to a health intervention, there will be a long transition period in which some of the labor force will have grown up in a poor health environment. This is important, because there is good evidence that many of the most important health interventions are those that affect young children (or even *in utero*). Specifically, children who grow up in a more favorable health environment are healthier in a number of measurable respects (such as height, IQ and prevalence of chronic disease), and perform better as students and workers.

To deal with this problem of phase-in, we allow a worker's health human capital to be a function of both the current health environment and the health environment that prevailed when he was born. In the case of the ASR measure, for example, human capital from health per worker of age i in period t, $h^h_{i,t}$, is computed as:

$$h^h_{i,t} = \begin{cases} \exp\left[\rho ASR\right] & \text{if } t-i < T-65 \\ \exp\left[\eta\rho ASR + (1-\eta)\,\rho ASR'\right] & \text{if } T > t-i \geq T-65 \\ \exp\left[\rho ASR'\right] & \text{if } t-i \geq T \end{cases} \tag{6}$$

where ASR and ASR' are the adult survival rates implied by the mortality regimes prevailing before and after the shock. $\rho > 0$ captures the birth-period health effects on lifetime productivity, and $\eta \in [0,1]$ captures the effects of an improvement in the health environment during an individual's working-age span. A value of $\eta = 1$ implies that health improvements are fully reflected in worker productivity right away. A value of $\eta = 0$ implies that there is no contemporaneous effect of health improvement on worker productivity; the only workers who will be more productive are those who are born after the improvement in health. At this point we have no solid grounds for estimating the value of η, and so in our simulations we consider values of 0 and 0.5. [In later drafts we will also consider the value of 1.0.]

To measure the effects of eradicating specific diseases, we use data on disease-specific YLDs and on disease specific mortality, both broken down by age. To apply this method we have to start with a specific life table. In the analysis below, we use the life table for Zambia in 2001, obtained from the WHO. (Zambia is fairly representative of sub-Saharan Africa as a whole. In 2001, its life expectancy at birth was 37 years. Malaria was the cause of about 8.3 per cent of deaths in Zambia, compared with a sub-Saharan average of 9.8 per cent. Tuberculosis was a little more severe than the sub-Saharan average, causing about 3.1 per cent of deaths compared with 2.0 per cent on average.) For each disease that we consider, we use data from the WHO on disease-specific deaths to create the corresponding cause-deleted life tables. We then use this cause-deleted life table to project the population from the shock-period (i.e., year 0) onward.[4] The effect of eliminating tuberculosis is to raise life expectancy at birth from 37.0 years to 38.0 years. Eliminating malaria would raise life expectancy at birth to 38.6 years.

[4] Strictly, we scale data on age-specific causes of death in sub-Saharan Africa by the population prevalence of that cause of death in Zambia compared with sub–Saharan Africa as a whole, since data on age-specific causes of death in Zambia are not available.

Our method for calculating the change in YLDs is to look at data for disease-specific YLDs broken down by age. We also have data on total YLDs by age. We simulate the eradication of a disease by subtracting disease-specific YLDs from total YLDs for each age group. For example, in the 45-49 age group, total YLDs per capita are .1403, and YLDs from tuberculosis are .0016. Thus, eliminating tuberculosis would lower YLDs per capita in this age group by a little over 1 per cent and raise labor input per capita by almost 0.2 per cent.

As with the effect of general improvements in health, we can consider different possibilities regarding the degree to which health shocks affect the productivity of those who are living at the time of the intervention compared with those who are born after it occurs. We assume that eradication of tuberculosis instantly eliminates all YLDs due to tuberculosis (so $\eta_{TB} = 1$), while some of the effects of malaria linger in the population that was previously exposed to it (so $\eta_{mal} = 0.5$). If we consider the elimination of several diseases at once (which we do not do in this draft), we can allow for different degrees of impact on living vs. future generations. Specifically,

$$h_{i,t}^{h} = \begin{cases} 1 - \sum_d \widetilde{YLD}_{d,i} & \text{if } t - i < T - 65 \\ 1 - \left\{ \sum_d (1 - \widetilde{\eta_d}) YLD_{d,i} + \sum_d \eta_d \widetilde{YLD}'_{d,i} \right\} & \text{if } T > t - i \geq T - 65 \\ 1 - \sum_d \widetilde{YLD}'_{d,i} & \text{if } t - i \geq T \end{cases}$$

where $\widetilde{YLD}_{d,i}$ and $\widetilde{YLD}'_{d,i}$ are the per-capita rates of YLDs from disease d affecting people of age group i in the pre and post-shock health regimes. η_d is a disease-specific coefficient indicating the extent to which cohorts alive at the time of the shock benefit from the shock.

2.4 Channels from Health to Income

Health may effect income through both demographic and economic channels. The simplest demographic channel involves changing the ratio of workers to dependents. A similarly simple channel works through experience. Changes in health may also affect education, savings rates and fertility rates.

The effect of an improvement in health on the age structure of the population is ambiguous. For example, the largest effect of many health improvements, such as reducing the prevalence of malaria, is to reduce the rate of child mortality. At least in the short run, this will raise the dependency ratio, movements in which in the developing world are dominated by changes in the ratio of children to those of working age. Naturally, a change in health which leads to relatively more dependents reduces per capita income, although not per worker income. However, in the long run, once fertility adjusts to the new survival rates, even an improvement such as reducing malaria may eventually reduce the dependency ratio through its effect on adult mortality.

Relatedly, an increase in adult survival rates implies that more of the workforce will be able to accumulate useful experience, thereby raising both per capita and per worker income. Again, though, to the extent that the initial effect of a health improvement is concentrated in younger

parts of the population, in the short-run it will reduce per-worker experience and therefore also income per-worker.

There are several possible channels through which changes in health may increase education. Longer life expectancy increases the time over which investment in human capital can be amortized, and therefore should raise investment in schooling. There is some evidence that healthier children are also better able to take advantage of schooling, for example through reduced absenteeism and greater mental alertness while at school (**?**). Another channel is the so-called "quality-quantity" trade-off. If disease eradication results in households having fewer children, the household budget constraint - and, at the macro level, the government's budget constraint - may be loosened, allowing greater investment in each child. At the moment, we take a reduced-form approach to evaluating the effect of health improvements on human capital in schooling, by allowing schooling to increase with some fixed elasticity in response to an improvement in life-expectancy.

Changes in health, particularly through changes in adult life-expectancy, may also cause changes in the savings rate. The classic life-cycle model of savings of Modigliani would suggest that an increased probability of surviving past the age of effective labor force participation would increase savings rates in the long-run. In the short-run, the demographic bulge of relatively young workers saving at a relatively high rate might also increase capital accumulation relatively shortly after the shock. Although these mechanisms may be important, it remains difficult to judge quantitatively how important life-cycle savings effects are likely to be in a developing economy. There is a lively discussion of the evidence on these issues as they relate to Taiwan in **?** and **?**.

3 Results

As already stressed, our interest is not in outlining the general pattern of co-movement of population measures and economic growth, since these co-movements arise in the data from the endogeneity of both demographic and economic variables. Rather, our concern is with identifying the economic effects of an exogenous change in health and the mortality regime. A natural way to communicate the results from our exercise is by showing projected paths of outcomes relative to a baseline in which the exogenous change in question does not take place. More specifically, our results include time series plots of variables in a base case and under alternative scenarios, or else of outcomes in particular scenarios relative to what would happen in the base case of no intervention.

As foreshadowed, we present results from two types of experiments, involving either an improvement in life-expectancy in general or an improvement relating to a specific disease. We consider first the effect of an improvement in life-expectancy at birth from $e_0 = 40$ to $e_0 = 60$. This approximately corresponds to the largest change in life expectancy induced by the health interventions of the 1930s and 1940s studied in **?**. The experiment employs model life table data for the South Asia region from the **?** and age-specific fertility data for Sri Lanka in 1953 from **?**. These data are among the earliest available and should capture much of the relevant demographic

behavior over the period considered by **?**. We consider two alternative measures of the effect of increased life-expectancy at birth on labor productivity, the ASR-based measure and the YLD-based measure.

Figure 2 shows the effect of the increase in life expectancy using the ASR-based measure of the productivity of health. Panel (a) shows the effect of the increase in life-expectancy on the size of the population, and breaks down the effect into the population of working age and the population of dependents. Since the increase in the dependent-age population exceeds the increase in the working-age population for the first forty-five years after the shock, the dependency ratio must increase over this period. This effect peaks about 15 years after the shock, at which time the dependency ratio has increased by about 0.10, from about 0.69 to about 0.79. Thereafter, the dependency ratio gradually declines to a long-run level of about 0.64. In the long-run, therefore, there is a demographic dividend in terms of income per capita from the decline in mortality, but this occurs only more than half a century after the shock.

The next three panels show three outcomes for three distinct cases, corresponding to different assumptions about the direct effect of improved health on labor productivity. The case labelled "Baseline" assumes that $\eta = 0.5$, so that cohorts alive at the time of the shock receive half of the productivity benefit, while cohorts born at the time of the shock or afterwards receive the whole of the productivity benefit. The case labelled "No Immediate Health" assumes that $\eta = 0$, so that only cohorts born at the time of the shock or after gain the labor productivity benefits of the shock. Finally, the case labelled "No Health" assumes that $\rho = 0$, that is, that there is no direct effect of health on labor productivity.

Panel (b) shows the evolution of physical capital per worker following the shock. Physical capital per worker falls following the shock mechanically because more workers are now alive to work with the same aggregate amount of capital. Since much of the increase in life expectancy is among the young, the cohorts entering the labor force after the shock are substantially larger than earlier incoming cohorts. This has the effect of depressing the capital stock per worker still further. Eventually, as the size of the population stabilizes, increased savings from the extra workers lead to a gradual recovery in the capital-labor ratio. The capital-labor ratio reaches a minimum about fifty years after the shock, at which point it is about 20 per cent lower than it was before the shock. Whether the capital-labor ratio eventually recovers to its former level depends on whether the increase in population and attendant increased pressure on the fixed factor, land, is so large as to outweigh the long-run positive effects (whether direct or indirect) of improved health on labor productivity. Since the positive effects of improved health are negligible in the "No-Health" case (as will be seen in panel (c)), the per capita capital stock never recovers in that case. In the cases where there is a direct positive effect of health on productivity, in the long-run the per worker capital stock increases by about 5 per cent. However, in either case the recovery is extremely slow; the supply of physical capital per worker only regains its pre-shock level about 120 years after the shock.

Panel (c) shows the evolution of human capital per worker. This is a combination of the direct effect of the health shock on labor productivity, and its effect through changing the age-structure, and therefore experience, of the workforce. The long-run effect of the changed age-structure of the workforce is to increase per worker human capital by about $4\frac{1}{2}$ per cent. The long-run effect of the shock on health human capital is of course identical regardless of whether existing cohorts receive its benefits or not, and is about 17 per cent.

Finally, panel (d) shows the overall effect of the shock on income per capita. Since in the long-run the dependency ratio falls and workers have more physical and human capital to work with, unsurprisingly income per capita rises. The long-run increase is about 10 per cent of income. The fact that income per capita falls in the "No Health" case indicates that the favorable effects of health on experience and the dependency ratio alone are insufficient to overcome the negative effect of increased population, working through greater pressure on the fixed factor. Whether existing cohorts receive the productivity benefits of health makes little difference to the horizon over which income per capita returns to its previous level, which is about fifty years in either case. But this assumption makes a significant difference to the intensity of the short-run fall in income per capita. In either case, income per capita troughs about 15 years after the shock. But in the baseline case, the fall in income per capita is only about 5 per cent of the previous level of income, whereas it is about twice that if the productivity of existing cohorts is not improved by the change in health conditions.

Figure 3 can be discussed more briefly than Figure 2, since it presents the same experiment, but using YLD-based measures of labor productivity rather than ASR-based measures. Of course, this has no effect on the projected path of population growth, and it turns out to have only a small effect on physical capital accumulation as well. But the changes in human capital are quite different. In general, YLD-based measures suggest smaller effects on effective labor supply than ASR-based measures. As a result, in the YLD experiment, human capital per worker actually falls marginally just after the shock, because the positive effects on productivity are more than offset by the shifting of the workforce towards younger ages. In the long-run, the YLD measure suggests that the health shock raises labor productivity directly by about 5 per cent, only about one-third of the estimate derived from the ASR-based measure.

Because the increase in labor productivity is so much smaller in the YLD-based measure of health human capital, the difference between the three scenarios is also smaller. Moreover, even in the long run, the small positive effects of health on productivity are insufficient to offset the deterioration in the land-labor ratio caused by the increase in population following the improvement in health. The YLD-based measure suggests that in the long-run, improvements in health are approximately neutral with respect to per capita income, as shown in panel (d).

The second experiment applies our theoretical framework to the dynamic effects of eradicating mortality from specific diseases. Specifically, we consider the alternative scenarios of completely eradicating the prevalence of malaria versus tuberculosis in Zambia, relative to the base case of

no intervention. The pre-shock mortality regime is generated from life table data for Zambia in 2001, obtained from the WHO. To simulate the appropriate shocks, we use data from the WHO on disease-specific deaths to create the corresponding cause-deleted life tables, which are then applied in the respective scenarios to project the population from the shock-period (i.e., year 0) onward. The direct effect of health on productivity is captured through YLD measures rather than ASR measures, since these seem more appropriate when considering the eradication of particular diseases.

As mentioned previously, Zambia is fairly representative of sub-Saharan Africa as a whole. Eradicating the diseases in question raises life expectancy at birth by only a year or a little more. Accordingly , despite the large number of deaths caused by these diseases, the effect of eradicating these diseases on life expectancy is small compared to the twenty years that characterized the international demographic transition in the previous experiment. Eliminating deaths from malaria increases life expectancy at birth to 38.6 years, while eliminating deaths from tuberculosis increases life expectancy at birth to just under 38 years. Since these improvements are much smaller than those considered in the first experiment, their economic effects will naturally also be substantially smaller.

The results of this experiment are reported in Figure 4. Panel (a) of Figure 4 presents the level of the population following the eradication of the diseases. Since malaria accounts for a greater fraction of mortality than tuberculosis, and this mortality is concentrated at younger ages, not surprisingly the increase in population resulting from its eradication is also larger: about 5 per cent in the long-run compared with about 2.7 per cent in the long-run. In both cases, around 80 per cent of the extra population growth occurs in the first forty years after the shock.

Panel (b) shows the most substantial difference in the effect of eliminating the two diseases. Eliminating malaria causes the dependency ratio to increase by about 2.6 per cent over the following 15 years, while eliminating tuberculosis causes the dependency ratio to fall more or less continuously for the next 60 years, including on impact. It is worth noting that the dependency ratio implied by the Zambian life tables is significantly higher before the shock than that implied by the South Asian model life tables in the earlier experiment. Before the shock, the dependency ratio is about 0.95 using these life tables, compared with only about 0.69 in the previous experiment.

Because of the differing effects on the size of the working-age population, the effects on physical capital per worker (panel (c)) are also rather different, with an immediate decline of about 2 per cent in the case of tuberculosis eradication followed by a gradual recovery to a level close to that before the shock. By comparison, although the initial decline in capital per worker caused by malaria is smaller (because malaria causes fewer working-age deaths), the eventual decline is about 50 per cent more pronounced and, even in the long run, capital per worker remains a little below its previous steady-state value.

The effect on human capital per worker also differs across diseases. Even the long-run effect on human capital of eradicating malaria is only about 0.2 per cent. This is because both malaria

fatality and morbidity are concentrated amongst those who are too young to work, so that the effect on the morbidity and age structure of the working population is small. By contrast, eradicating tuberculosis reduces mortality and morbidity mainly amongst prime age workers, thereby skewing the age distribution of the population towards relatively more experienced workers. As a result, the human capital effect per worker of eradicating tuberculosis is about 4 times as great as that of eradicating malaria.

The net effect of these various differences is that eradicating tuberculosis has a slight positive impact on GDP per capita at the time of eradication, which climbs over time to about a 1 per cent increase in the long-run. By contrast, malaria eradication reduces GDP per capita by about 1.6 per cent at a 35 year horizon, and has no discernible long-run effect on GDP per capita.

4 Sensitivity Analysis

This section presents several sensitivity analyses based on varying parameter values in the model. We focus on how the results change relative to the baseline case in the experiment involving a shock to life expectancy and an ASR-based approach to health human capital. The three exercises that we perform involve the role of experience, the rate of fertility adjustment and the share of land in national income.

Figure 5 highlights the role of experience in the model. Accounting for the human capital acquired through experience somewhat increases the amplitude of the changes following the shock. Human capital per worker (panel (a)) grows less during the period from 15 to 50 years after the shock, during which time the younger cohorts are relatively large. After that, however, the long-run shift in the age-structure of the population is towards slightly older workers. This implies that, as can be seen in panel (b), income per capita falls less by about one per cent or so in the case without experience than in the case with experience, while the long-run gain in income is about 1.7 percentage points greater when accounting for the role of experience.

Figure 6 shows how the results are changed by changing the horizon over which fertility adjusts. The baseline case assumed that it would take 50 years for fertility to adjust to its new long-run rate. The figure also shows fertility adjustments that take 25 years, which is perhaps a little closer to the historical experience of fertility adjustment in East Asia, or 75 years, which is perhaps closer to the historical experience of Europe or the experience of some sub-Saharan African countries in the more recent past.

Panel (a) shows the long-run effect on population size. Relative to no change in life-expectancy, the population increases by about 31 per cent, 52 per cent and 76 per cent respectively as fertility takes longer to adjust. Differences in the rate of fertility adjustment become apparent only fairly gradually. After 25 years, the population is 20 per cent, 24 per cent and 26 per cent bigger in the three scenarios. After 50 years, however, the differences are apparent, with the population increase only 27 per cent in the rapid adjustment case, but 42 per cent in the baseline

case and over 50 per cent in the slow adjustment case. Panel (b) shows that, not surprisingly, slower adjustment of fertility exaggerates the short-run fall in income per capita and reduces the long-run increase. In all three cases, income per capita falls by about 5 per cent 15 years after the shock, but intriguingly, in all three cases, income per capita recovers to close to its original level in the same amount of time that fertility takes to adjust.

Finally, Figure 7 shows how the results of the model are altered when the income share of land is increased from 10 per cent to 20 per cent, and the capital share is correspondingly decreased. Our baseline case, in which the land share is 10 per cent, is probably quite conservative for most developing countries. In a well known study, **?** assume a value of 30 per cent for pre-industrial economies. Evidence from **?** also suggests that a land share of 10 per cent is likely to be conservative. In Zambia, for instance, the World Bank reports two measures of the importance of fixed factors. First, it reports that natural capital (the value of pasture and cropland, subsoil resources, etc.) equaled $1779 per capita in the year 2000, as compared to $694 for produced capital. The same source reports that resource rents in that year constituted 12% of Zambia's GDP. Accordingly, a sensitivity analysis that increases land's share of national income to 20 per cent seems reasonable for at least some developing countries.

Figure 7 shows that there is almost no apparent difference between the two simulations in GDP per capita in the first half century after the shock, and indeed that income recovers to its pre-shock level at almost the same time. They have, however, quite different implications for the long-run gains from improvements in health. Doubling the share of land more than halves the long-run gains in per capita income from improved health, which go from 10 per cent to 3.5 per cent. Naturally, the longer fertility takes to adjust, the more pronounced this effect will be.

5 Concluding Remarks

Using a simulation model, we explore a number of different dynamic effects of an exogenous intervention in a population's health. Our initial analyses show that for reasonable parameters, the phase-in period can be very long, on the order of half a century, before any beneficial effects of an improvement in health are visible in GDP per capita. It may take twice that long to achieve most of the long-run gains in income per-capita resulting from increased health. This lag in the effect of health on income per capita may explain one of the more puzzling phenomena regarding cross-country inequality. While cross-country inequality in health declined rapidly over the period 1950-1990 (that is, up through the advent of the AIDS epidemic), the level of cross-country inequality in income did not.

The results from our analyses of health's effect on economic growth will have a number of uses. Considerations of economic effects are already an important part of discussions of and advocacy for programs to improve population health. While health improvements may well raise worker productivity, many potential interventions in developing countries will also be accompanied

by the side effect of a rapidly growing population, which will have negative economic effects over a significant time horizon. An understanding of the demographic dynamics that accompany health improvements may therefore suggest complementary investments, for example in family planning, which can greatly improve the rate at which improvements in health are translated into improvements in the standard of living.

It is appropriate, though, to end on a note of caution. That improvements in health may temporarily (or even permanently) reduce income per capita is not, of course, a reason not to pursue such improvements, which are valuable in themselves. Similarly, family-planning policies need to be considered in the context of welfare analysis rather than simply through the lens of their effects on per-capita income. This study is therefore complementary to the consideration of the welfare analysis of development policies, not a substitute for it.

References

Acemoglu, Daron and Simon Johnson, "Disease and Development: The Effect of Life Expectancy on Economic Growth," 2006. Working Paper, Department of Economics, MIT.

Barro, Robert J. and Jong-Wha Lee, "International Data on Educational Attainment: Updates and Implications," 2000. Working Paper, NBER.

Bils, Mark and Peter J. Klenow, "Does Schooling Cause Growth?," *American Economic Review*, December 2000, *90* (5), 1160–1183.

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

Commission on Macroeconomics and Health, *Macroeconomics and Health: Investing in Health for Economic Development*, Geneva: World Health Organization, 2001.

Deaton, Angus, "Saving and Growth," in Klaus Schmidt-Hebbel and Luis Servén, eds., *The Economics of Saving and Growth: Theory, Evidence, and Implications for Policy*, New York, NY: Cambridge University Press, 1999, pp. 33–70.

Hall, Robert E. and Charles I. Jones, "Why Do Some Countries Produce So Much More Output Per Worker Than Others?," *Quarterly Journal of Economics*, February 1999, *114* (1), 83–116.

Hansen, Gary D. and Edward C. Prescott, "Malthus to Solow," *American Economic Review*, September 2002, *92* (4), 1205–1217.

Keyfitz, Nathan and Wilhelm Flieger, *World Population: An Analysis of Vital Data*, Chicago, IL: University of Chicago Press, 1968.

Lee, Ronald D., Andrew Mason, and Tim Miller, "Saving, Wealth and Population," in Nancy Birdsall, Allen C. Kelley, and Steven W. Sinding, eds., *Population Matters: Demographic Change, Economic Growth, and Poverty in the Developing World*, New York, NY: Oxford University Press, 2001, pp. 137–164.

Mills, Anne and Sam Shillcutt, "Communicable Diseases," in Bj
orn Lomborg, ed., *Global Crises, Global Solutions*, New York, NY: Cambridge University Press, 2004, pp. 62–114.

United Nations, *Model Life Tables for Developing Countries*, New York, NY: United Nations Population Division, 1982.

Weil, David N., "Accounting for the Effect of Health on Economic Growth," 2007. Forthcoming, Quarterly Journal of Economics.

World Bank, *Where is the Wealth of Nations? Measuring Capital for the 21st Century*, Washington, DC: World Bank, 2005.

Young, Alwyn, "The Gift of the Dying: The Tragedy of AIDS and the Welfare of Future African Generations," *Quarterly Journal of Economics*, May 2005, *120* (2), 423–466.

Table 1

Population and YLDs for Males Aged 30-44, AFRO E Region		
	Total	Per Capita
Population	26,002,500	N/A
All Causes	3,514,640	0.135
Communicable, Maternal, Perinatal and Nutritional Conditions	1,055,733	0.041
Infectious and Parasitic Diseases	1,032,689	0.040
Tuberculosis	119,800	0.005
HIV/AIDS	535,967	0.021
Malaria	35,114	0.001
Non-communicable Diseases	1,865,649	0.072
Injuries	593,258	0.023

Figure 1: YLD and Life Expectancy

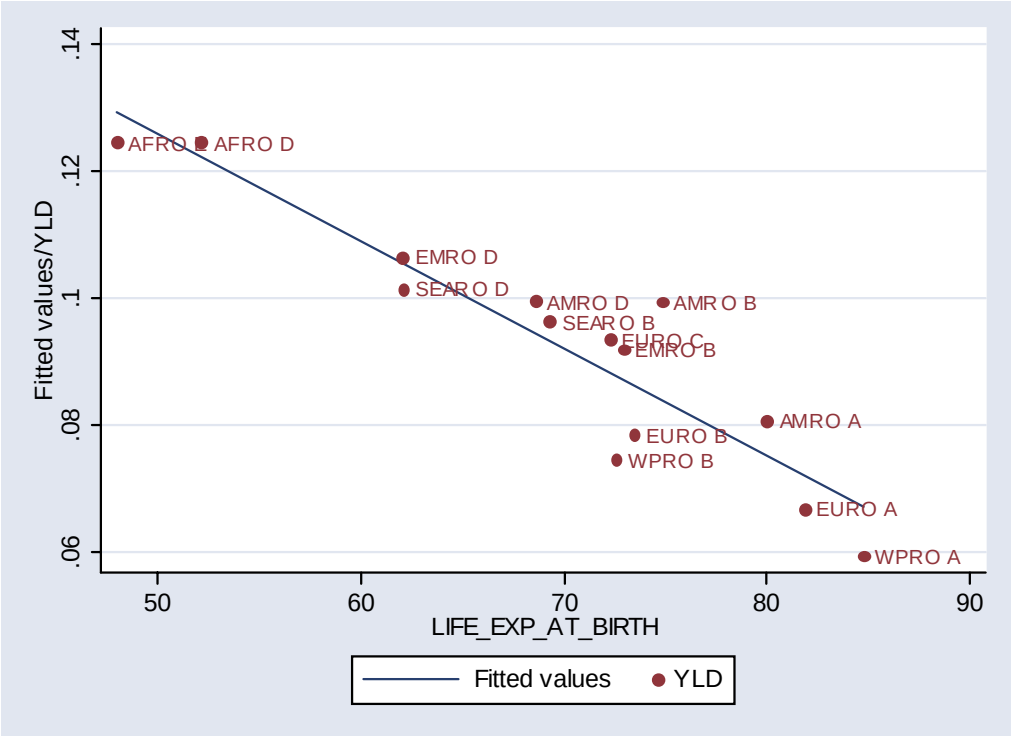

Figure 2: Shock to Life Expectancy (Health: ASR)

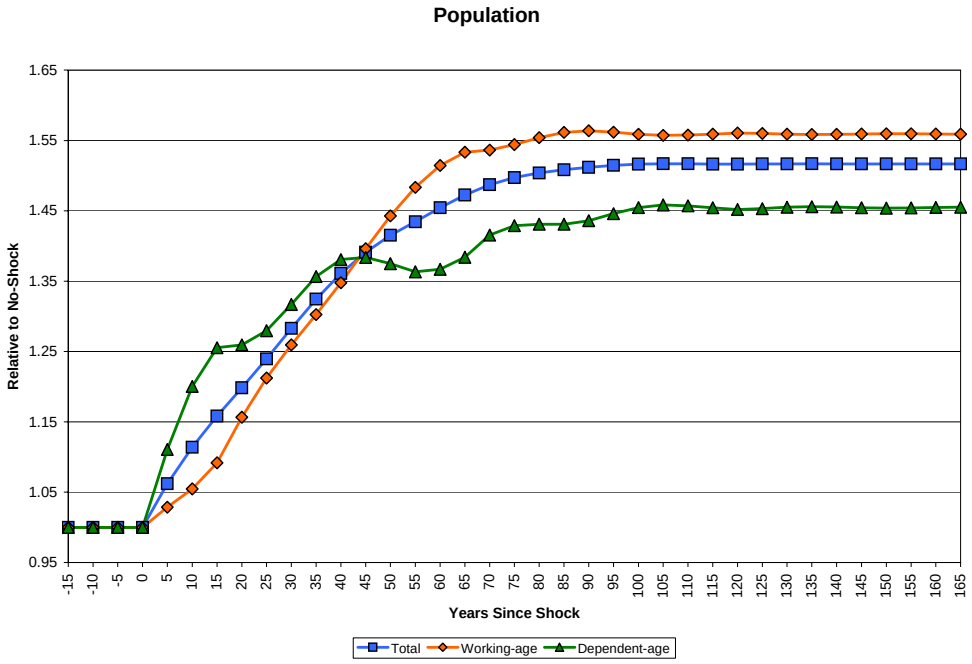

Panel (a)

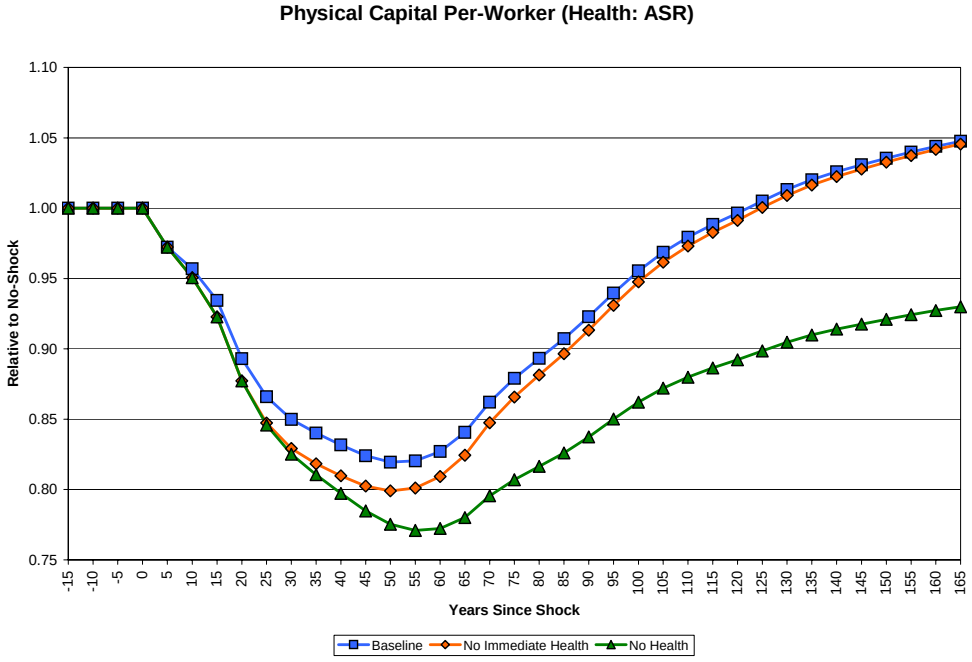

Panel (b)

Figure 2: Shock to Life Expectancy (Health: ASR)

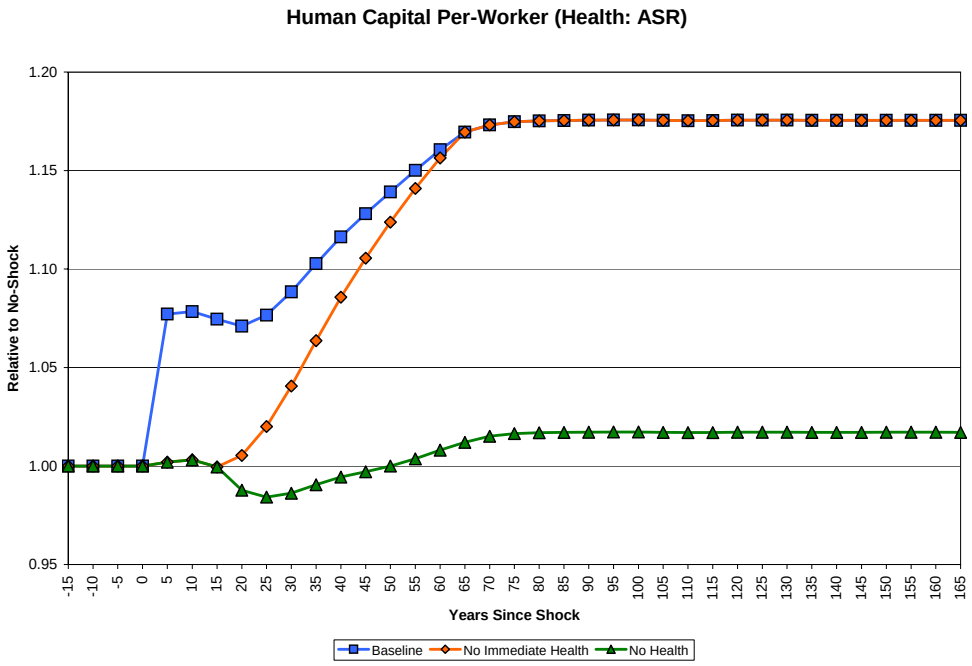

Panel (c)

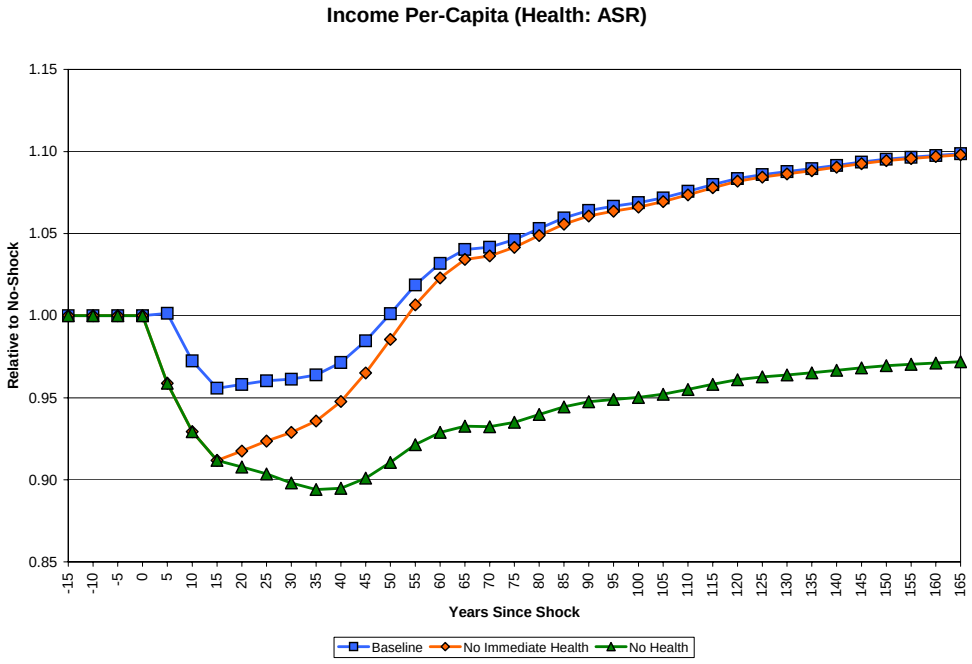

Panel (d)

Figure 3: Shock to Life Expectancy (Health: YLD)

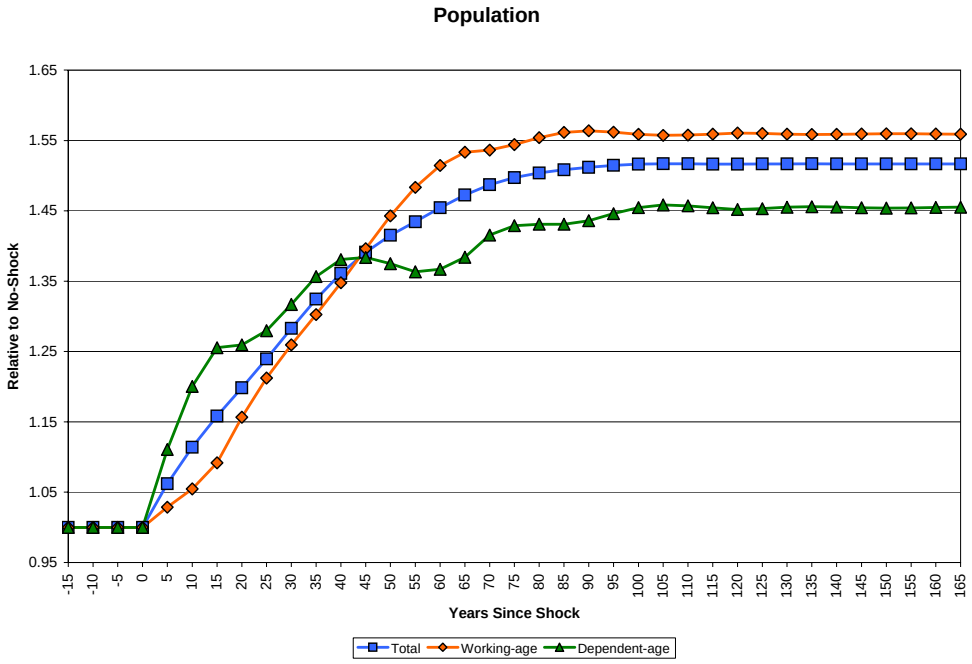

Panel (a)

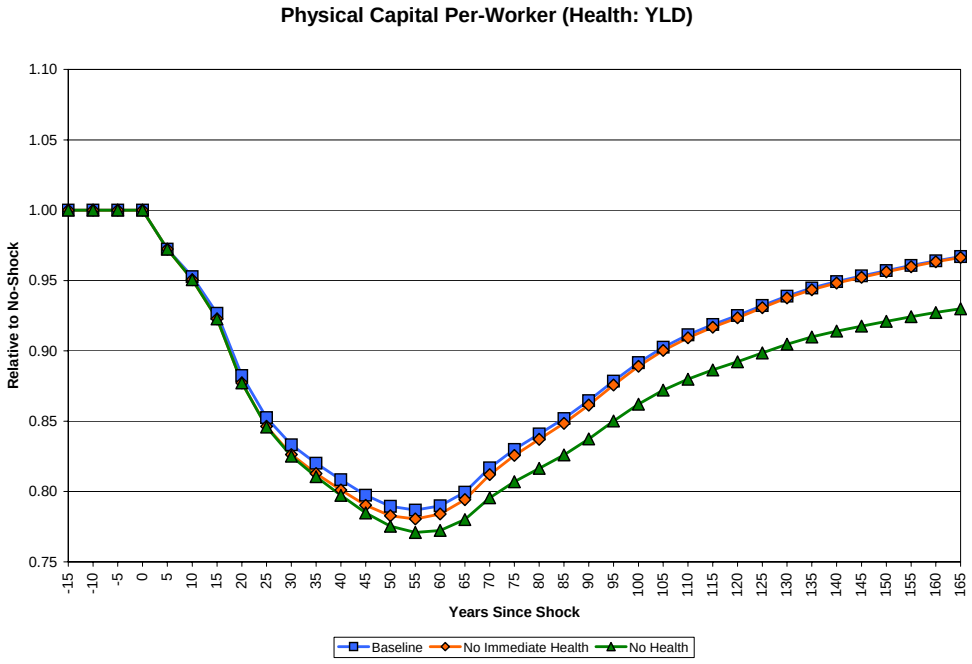

Panel (b)

Figure 3: Shock to Life Expectancy (Health: YLD)

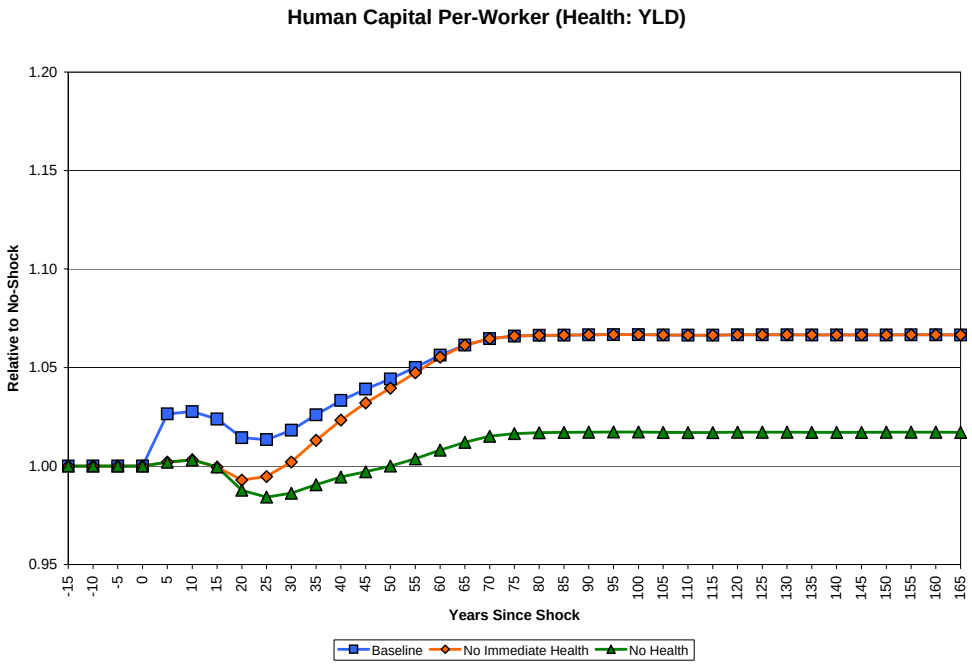

Panel (c)

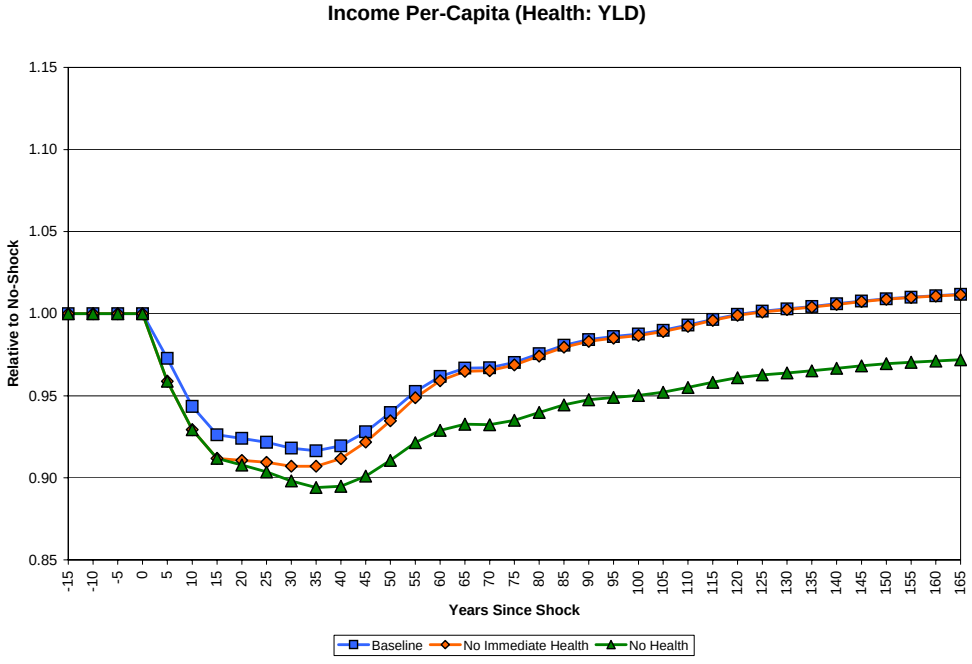

Panel (d)

Figure 4: Disease Eradication

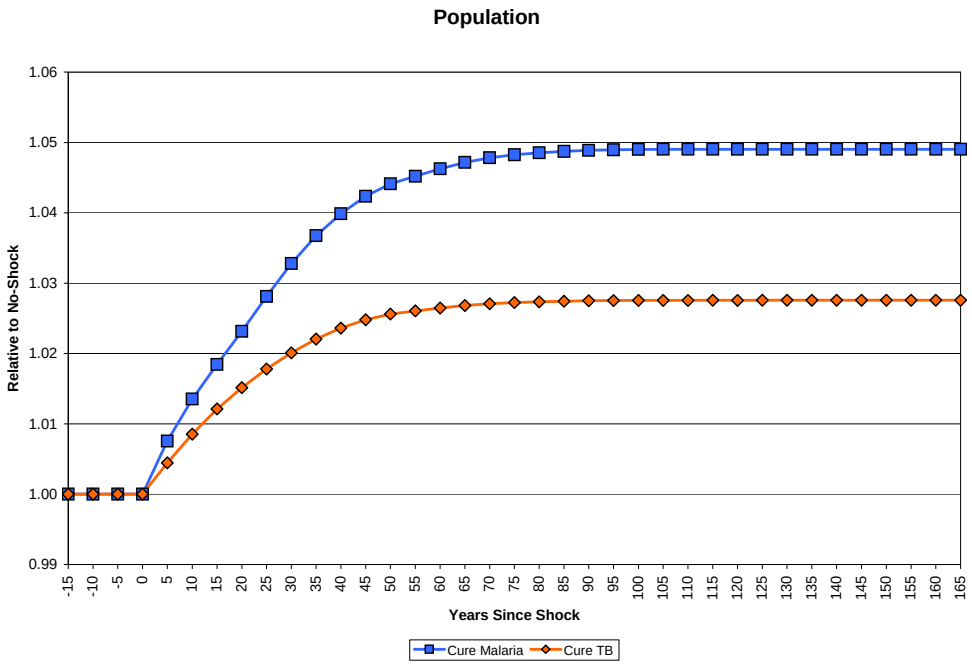

Panel (a)

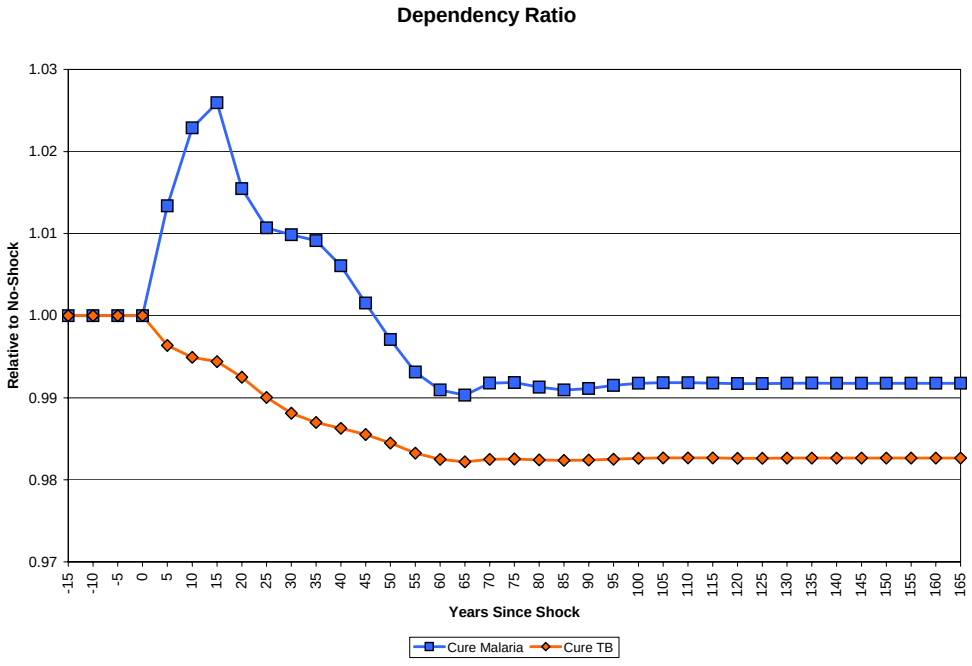

Panel (b)

Figure 4: Disease Eradication

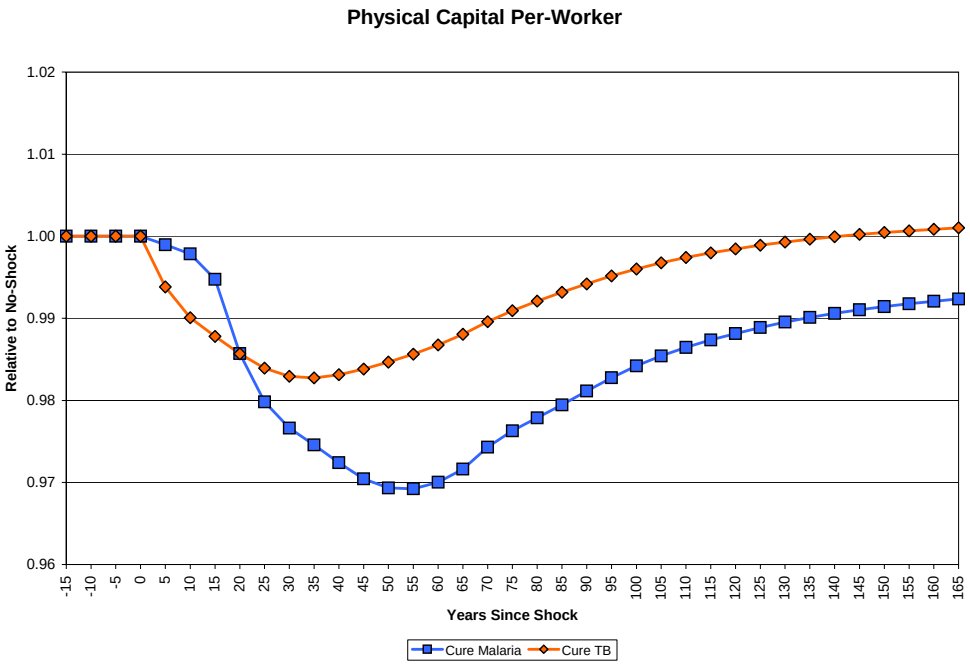

Panel (c)

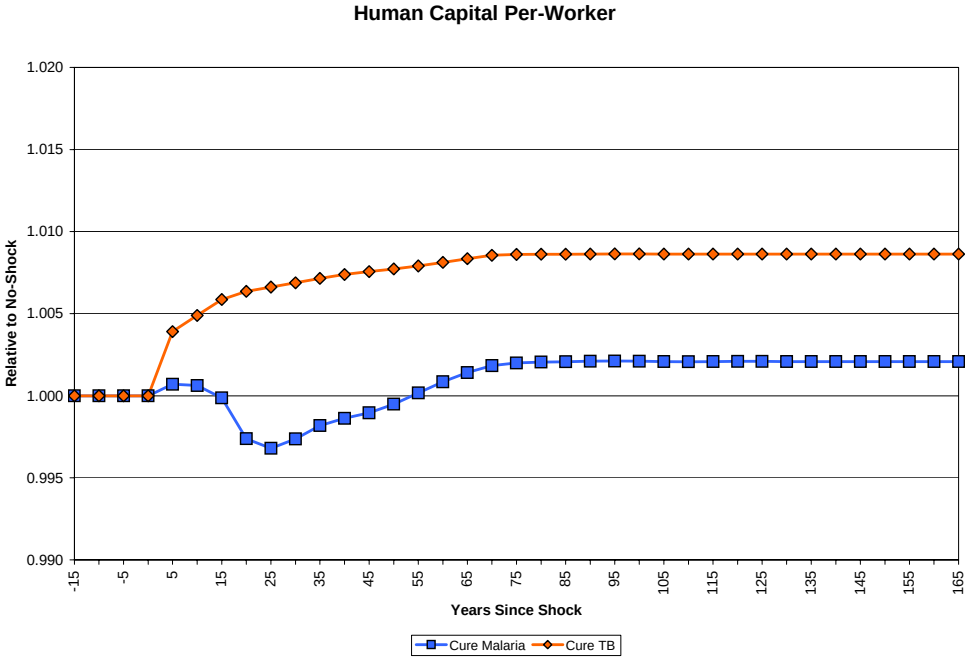

Panel (d)

Figure 4: Disease Eradication

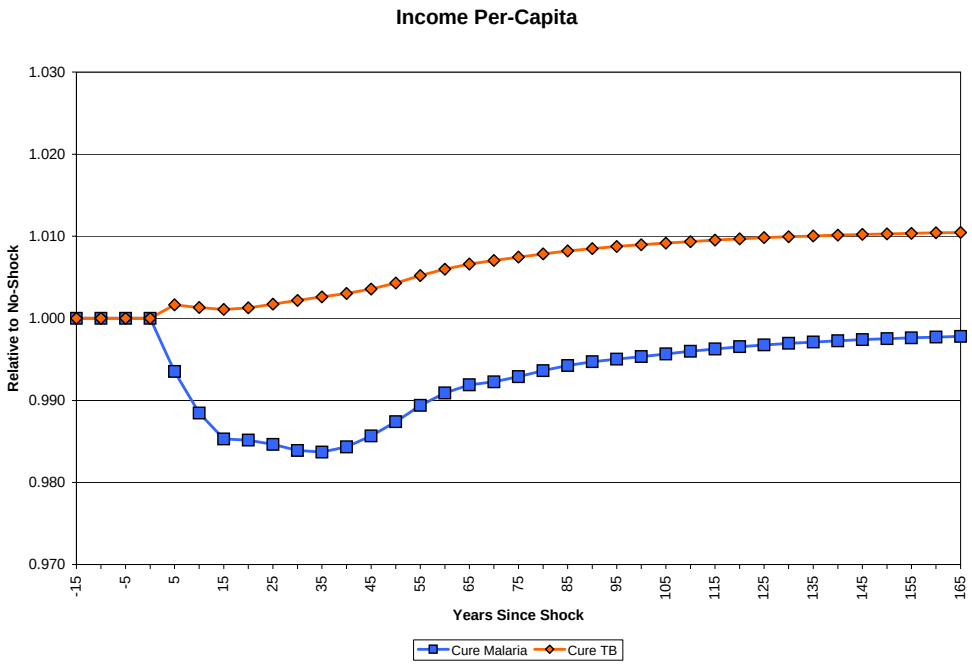

Panel (e)

Figure 5: Sensitivity to Experience Human Capital

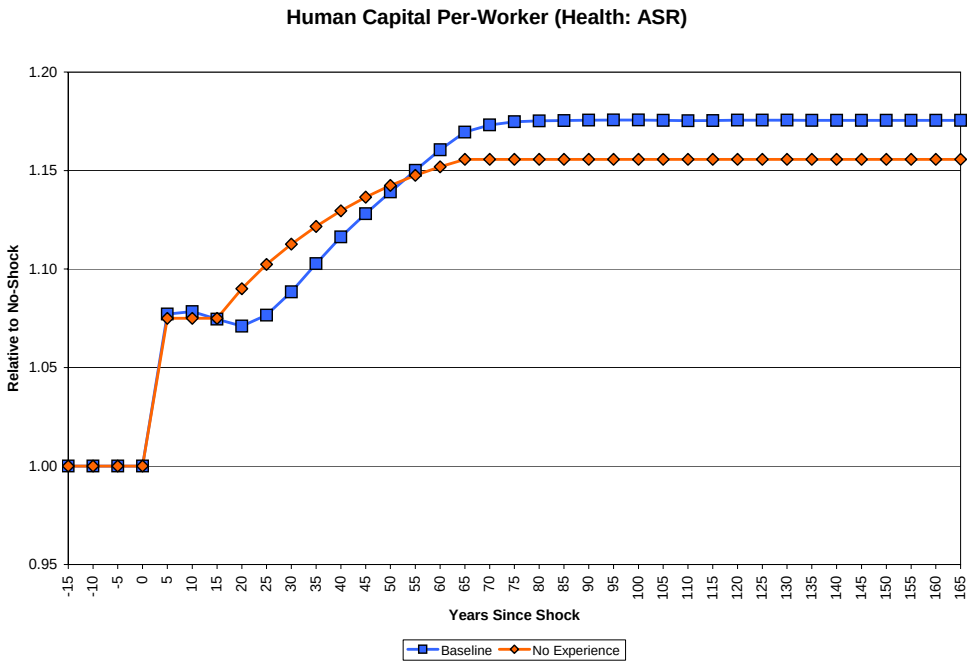

Panel (a)

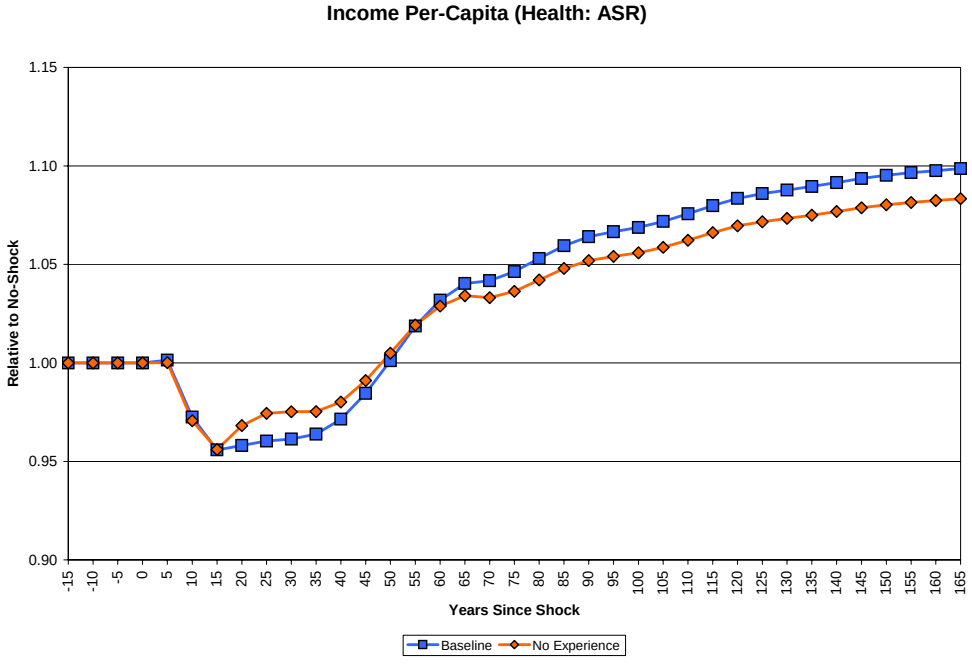

Panel (b)

Figure 6: Sensitivity to Fertility Adjustment Horizon

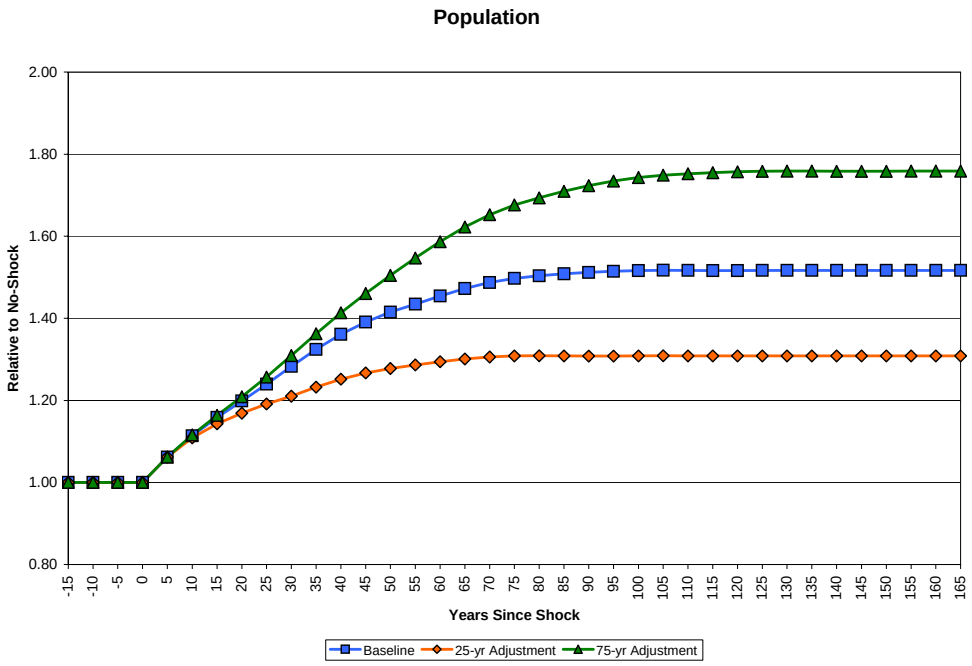

Panel (a)

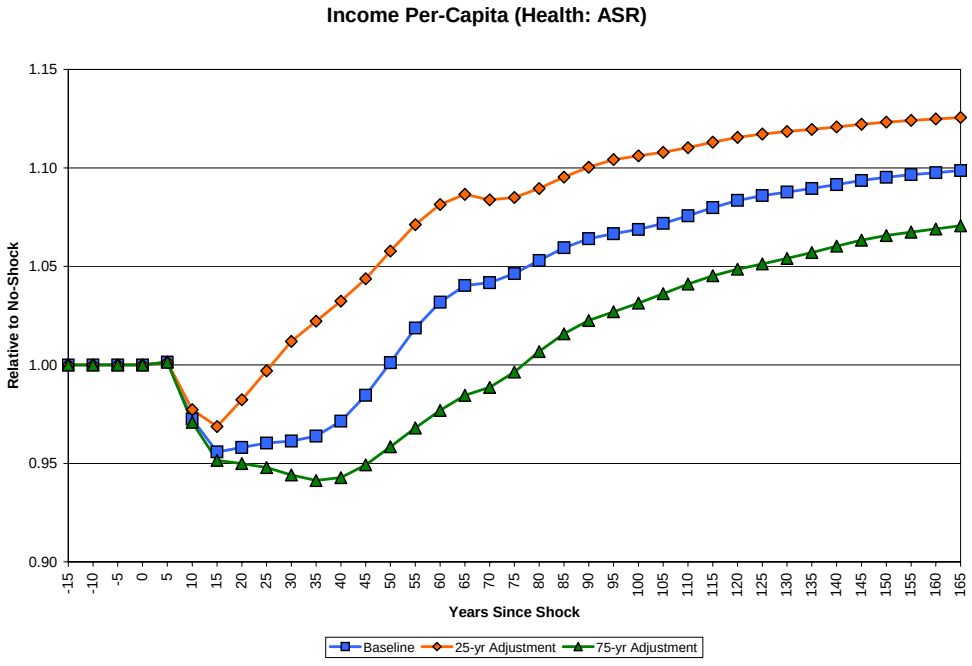

Panel (b)

Figure 7: Sensitivity to Land Share of Income