University of Connecticut

Department of Economics Working Paper Series

Malaria: Disease Impacts and Long-Run Income Differences

Douglas Gollin
Williams College

Christian Zimmermann
University of Connecticut

Working Paper 2007-30R

August 2007, revised April 2010

341 Mansfield Road, Unit 1063
Storrs, CT 06269–1063
Phone: (860) 486–3022
Fax: (860) 486–4463
http://www.econ.uconn.edu/

This working paper is indexed on RePEc, http://repec.org/

Abstract

Malaria is a parasitic disease that causes over 300 million ”acute illness” episodes and one million deaths annually. Most occur in the tropics, especially sub-Saharan Africa. Countries with high rates of malaria prevalence are generally poor, and some researchers have suggested a direct link from malaria to poverty. We explore the interactions between malaria and national income, using a dynamic general equilibrium framework with epidemiological features. We find that without prevention or control, malaria can have a large impact on income. However, if people have any effective ways of avoiding infection, the disease has little effect on income levels.

Journal of Economic Literature Classification: I1, O11, E13, E21

Keywords: Malaria, Epidemiology, GDP, Disease prevention, Sub-Saharan Africa.

Portions of this research were undertaken while Gollin was on leave at the Economic Growth Center, Yale University, and at the Centre for Study of African Economies, Oxford University. Gollin gratefully acknowledges support and facilities at both institutions. Hoyt Bleakley has kindly shared information and preliminary results. We have also benefited from the comments of Ben Bridgman, Steve Meardon, Lara Shore-Sheppard, Gustavo Ventura, David Weil, and seminar and conference participants at the Arizona State University Conference on Economic Development; the Society for Economic Dynamics 2005 meetings in Budapest; the University of Essex; Rutgers University; the University of Delaware; University of Connecticut; University of Houston; Wesleyan University; LAMES 2006 in Mexico City; the Northeast Universities Development Conference 2005 at Brown University; and the Harvard Center for International Development.s May 2007 conference on Health Improvements for Economic Growth.

1 Introduction

The World Health Organization (WHO) reports that malaria, a parasitic disease transmitted by mosquitoes, causes over 300 million episodes of "acute illness" and more than one million deaths annually.[1] Most of the deaths occur in poor countries of the tropics, and about 90 percent occur in sub-Saharan Africa. Infants and children account for most of the mortality from malaria; the disease is thought to account for one of every five child deaths in the world.[2]

Most of malaria's ravages are concentrated in countries that are very poor. Even within countries, malaria disproportionately affects poor people. According to the United Nations Children's Fund (UNICEF), "Malaria is truly a disease of poverty. It afflicts primarily the poor, who tend to live in malaria-prone areas in dwellings that offer few, if any, barriers against mosquitoes" (UNICEF 2005). Sachs and Malaney (2002) argue that "[a]s a general rule of thumb, where malaria prospers most, human societies have prospered least.... The extent of the correlation suggests that malaria and poverty are intimately related."

The causality of this relationship is complicated, however. Does malaria cause poverty? Or does poverty cause malaria? Both channels of causation seem reasonable. It is also possible, as noted by Sachs and Malaney (2002), that the correlation could be spurious, caused perhaps by some other direct connection between climate and geography with growth rates or income levels. Resolving these causality issues has been difficult for researchers trying to assess the economic impact of malaria.

The point matters; if malaria plays a significant role in reducing income per capita, then an intensive campaign of eradication or prevention might have large economic benefits – affecting not only individuals' well-being but also having impact at a macro level. Alternatively, if there is no strong relationship between malaria and national income, then efforts to control malaria should be based entirely on their welfare impacts: their effectiveness in reducing mortality, (especially infant and child deaths) and morbidity. Moreover, if the causal connection flows from income to illness, then policies of income support might be the most promising way to eradicate malaria or

[1] Reported by WHO on the "Roll Back Malaria" program website at: http://mosquito.who.int/cmc_upload/0/000/015/372/RBMInfosheet_1.htm, January 30, 2005.

[2] Reported by the United Nations Children's Fund (UNICEF), http://www.unicef.org/health/index_malaria.html, accessed June 10, 2005.

reduce its consequences.

Given the large numbers of deaths and the high toll that malaria takes on the world's poorest people, it seems likely that prevention and control of malaria, even if only partly successful, should cause large welfare gains. This paper does not question that view, nor does it ask whether malaria prevention and control measures are worthwhile, from an economic perspective. This paper asks the narrower question of whether malaria prevention and control measures, if successful, would be likely to increase average per capita income in countries with malaria burdens.

Our paper differs from previous efforts in several respects. We recognize that people in malaria-endemic economies have access to a number of potential behavioral responses that may mitigate their exposure to the disease or the impact of the disease. They may, for example, purchase bednets or mosquito sprays; they may purchase drugs to treat infants and children who have contracted malaria; they might even move to locations with lower malaria exposure. These responses are often omitted from studies of the disease's impact.

Our paper shows that a plausibly calibrated model can yield large impacts for a disease like malaria – though not so large as to account for much of the cross-country dispersion in income per capita. We further show that where behavioral responses to the disease are possible, the economic impact of the disease is likely to be diminished to a large extent. Where people have access to effective means of prevention or cure, we find that they face strong incentives to use them – even if these measures are quite costly. In our model, the extent of the behavioral response is limited by efficacy considerations, more than by the price of prevention tools. This runs counter to many current malaria policies (such as subsidized distribution of bednets). Our results suggest that re-focusing malaria policy on improvements in efficacy, rather than on direct subsidies, may have a positive impact.

This paper uses a heterogeneous agent model of an economy in which individuals face exposure to a malaria-like disease. The model disease is highly simplified, but it captures a number of key features of malaria. Specifically, our model includes the following features:

- People face a chance of being exposed to the disease in every period of their lives. If exposed, they face a probability of infection.

- People who become sick face reduced effectiveness as workers and increased probabilities of dying.
- An externality characterizes disease dynamics: specifically, the probability of becoming infected in any particular period depends on the proportion of the population affected.
- Our model economy features flexible factors and employs a constant returns aggregate production technology.
- We assume that there are some costly methods of prevention available, which may (or may not) be entirely effective. These might correspond to drugs, vaccines (such as those currently under development), or any of a number of existing methods of mosquito control, such as bednets or sprays.

At the same time, our model abstracts from the specifics of the disease in some important ways. Perhaps the most significant are the following:

- We do not specifically model a vector of transmission; i.e., there are no mosquitoes in our model. We can, however, proxy for differences in mosquito density and virulence.
- We do not model the transmission from individual to individual, so we are not concerned with specific chains or networks of infection. This would be important for a disease like HIV/AIDS, but it is less important for malaria.
- We do not explicitly model fertility behavior in the model. Instead, we impose some mechanical assumptions such that fertility rates will be highest in those countries with high disease burdens, as is observed in the data.

Our approach differs from previous research in several respects. First, we are particularly interested in the general equilibrium effects of the disease through its prevalence and factor prices, and through the behavioral responses that show up in general equilibrium. Second, we use a heterogeneous agent framework in which different individuals experience different health outcomes and in which malaria influences savings and capital accumulation. Third, we approach the problem with a different methodology than the existing literature, which is largely empirical. We argue that a theoretical

model, with appropriate calibration, can teach us important lessons in a subject area characterized by limited data and poor measurement. The empirical literature focuses on issues such as the uptake and use of bednets when they are given to households (or sold at varying prices). We prefer to focus on the underlying private demand for bednets. Moreover, because of heterogeneity and externalities, we prefer to focus on aggregate impacts, rather than individual impacts.

Our analysis finds that countries facing a severe malaria burden may in fact face large reductions in income. A country in which malaria is widely preventable and essentially uncontrollable may find its income reduced by almost half, relative to a situation in which the disease could be eradicated.

However, our analysis also suggests that the impact of the disease on income per capita falls sharply if there are any effective measures that allow individuals to prevent, control, or cure the disease. Because the private benefits of preventing malaria are very high, we find that people are willing to spend large fractions of their income to reduce their risk of contracting the disease – and more so as the severity of the disease (or its economic consequences) rises. As a result, in a world with costly but effective methods of prevention and control, we find that the disease will have relatively minor impacts on per capita incomes (though potentially larger impacts on welfare). This is important in terms of the implied policy prescription. Reducing the cost of prevention for users leads to minor improvements in our model; increasing the efficacy of prevention methods has large impacts.

The remainder of this paper is organized as follows. Section 2 provides some background and discusses previous literature on the subject; Section 3 presents our model; Section 4 describes the calibration. Section 5 presents the results of a quantitative exploration of the model, and Section 6 concludes.

2 Background

Malaria is an ancient disease, although its exact origins and evolutionary history are unclear. It was described in China some five thousand years ago. It is thought to have originated in Africa and to have spread subsequently into Asia and the Mediterranean. Greek writers recognized the disease and its symptoms, and one source notes that malaria was responsible for the decline of city-state populations and depopulation of

rural areas.[3] The disease appears to have migrated to the New World following the Columbian exchange, and to this day, fewer different strains of malaria are found in the Americas than in Africa and Asia. Recent research suggests that the origins and spread of the disease in the Old World paralleled the spread of sedentary agriculture (Tishkoff et al., 2001).[4]

Historically, malaria was endemic in most regions of the world. The morbidity and mortality burden of malaria differ from country to country, in part because the prevalence of the disease (and the conditions that give rise to the disease) differ substantially across regions. Hamoudi and Sachs (1999) report that historically, malaria was found as far north as 64 ° N latitude (farther north than Stockholm or Moscow) and as far south as 32 ° S.

2.1 Disease biology and ecology

Malaria is caused by a family of macroparasites that infect humans. There are in fact four species of *Plasmodium* parasites that cause malaria in people. These four species have similar life cycles; all are transmitted to humans by a mosquito vector (various species of *Anopheles* mosquitoes) and live a portion of their life cycle in the mosquito host.

A person is infected with malaria when he or she is bitten by an infected mosquito, which passes the *Plasmodium* parasite into the person's bloodstream in a form known as a sporozoite. The parasites lead a complex life cycle inside the human host, living at various stages in liver cells and red blood cells. From time to time, they cycle through stages in which they destroy numerous red blood cells. It is at this stage that the disease generates its most severe symptoms in infected people. Eventually, the parasites become gametocytes which are in turn ingested by mosquitoes that bite the human host. Inside the mosquito, the gametocytes mature, reproduce sexually, and migrate into the mosquito's salivary glands, at which stage the life cycle is repeated.[5] For some species of *Plasmodium,* the parasites may persist in the liver for months or years, resulting in chronic and recurring eruptions of merozoites that correspond to

[3] See http://www.cdc.gov/malaria/history (cited June 2005).

[4] See also McNeill (1976), pp. 219-221.

[5] The life cycle is described and illustrated at: http://www.dpd.cdc.gov/dpdx/HTML/Malaria.htm (viewed June 2005).

episodes of fever and sickness.

The disease varies with the infecting species of *Plasmodium* and with the individual's prior health and immune status. Typically, it causes fever and chills, along with headaches, vomiting, and diarrhea. It may also cause long-term anemia, liver damage, and neurological damage. The most dangerous species, *P. falciparum*, can cause cerebral malaria, a frequently fatal condition involving the brain and central nervous system. Those who survive cerebral malaria may experience lasting brain damage.

The prevalence of the disease varies across the globe, largely due to differences in the human exposure to *Anopheles* mosquito bites. Some of this variation is geographic and climatic: these mosquitoes are not found in areas of intense cold or in deserts (Sachs and Malaney 2002). Human exposures are also reduced in areas where mosquitoes spend winter months as eggs or in dormant stages of their life cycle. Exposures may also be reduced in areas where people spend significant fractions of their time indoors in enclosed or screened buildings, or where people are dressed in ways that will reduce exposure.

The ecological adaptation of different mosquito species is also important. Although many species of *Anopheles* mosquitoes are capable of transmitting the *Plasmodium* organisms, transmission occurs only when a mosquito first bites an infected human and then subsequently bites another (uninfected) human. Some species of mosquitoes, however, prefer not to feed on humans (although they will do so if other food sources are not available). Others are anthropophilic; *i.e.*, they prefer to feed on humans. Anthropophilic mosquitoes are obviously more likely to transmit malaria from individual to individual. Thus, areas where anthropophilic mosquitoes are prevalent are likely to face more acute malaria burdens.

McNeill (1976) notes that the geographic distribution of mosquito species is largely due to chance, from the perspective of humans. The distribution depends on highly local ecological differences (trace minerals in the water, salinity of water, types of habitat, etc.). Thus, pure ecological chance had large effects on the relative prevalence of different mosquito species, across the globe, and hence on the relative prevalence of malaria. McNeill notes that "the mosquito species which is Europe's most efficient vector of malaria... prefers to feed on cattle. If enough alternate sources of blood are available to them, these mosquitoes will eschew potential human hosts and thus

interrupt the chain of infection, since cattle do not suffer from malaria" (p. 117).

In fact, one of the apparent reasons for the extensive malaria burden in sub-Saharan Africa is the prevalence of two species of highly anthropophilic species of *Anopheles* mosquitoes: *An. gambiae* and *An. funestus*. These two species together inhabit much of the humid zone of Africa, and only the northern and southern extremities of the continent are free from these strongly anthropophilic species.[6] This clearly plays a significant role in accounting for malaria's impact in the region, although poverty may also play an important part (Sachs and Malaney 2002).

It is true that the types of mosquito prevalent in different regions are, in an ecological sense, related to human impacts on the landscape. It is also true, however, that the distribution of mosquito species across the landscape is largely exogenous in the short run. For the purposes of this paper, we will treat mosquito habitats as an exogenous characteristic of a place.

2.2 Economic impact of malaria

Malaria's most obvious economic impacts occur from mortality and from the loss of time associated with episodes of morbidity. Even when people survive malaria, however, the disease can occasionally cause lasting health and cognitive problems. It is associated with maternal anemia during pregnancy, with low birth weight for babies, and it is a major cause of childhood anemia. Severe disease episodes (*i.e.*, "cerebral" malaria) have been shown to cause severe long-term physical and neurological disability. There is no clear evidence on the cognitive impact of malaria on individuals who contract less severe cases of the disease, although there are some reports of non-trivial effects on learning among schoolchildren.[7]

There is at present no effective vaccine or inoculation to prevent malaria.[8] However, the disease can be treated at relatively low cost (at least in its milder forms) with drugs or even simple measures to reduce the severity of symptoms. Prevention measures are also relatively inexpensive. For example, mosquito nets impregnated

[6]The distribution of *Anopheles* species around the globe is shown at: http://www.cdc.gov/malaria/biology/mosquito/map.htm.

[7]A useful survey is Holding and Snow (2001).

[8]This is, however, an active area of medical research, heavily funded by international governmental and non-governmental donors.

with insecticides, available for $5-$10 each (or less), can significantly reduce exposure to mosquitoes and thereby limit malaria morbidity and mortality.[9] Intensive educational campaigns and project-level interventions have increased the use of bednets in Africa, with coverage estimated at 50-75 percent in many countries. But rapid growth in bed net coverage does not yet seem to have generated comparable declines in malaria deaths – although data quality is admittedly poor.

In assessing the overall economic impact of the disease, Sachs and Malaney (2002) survey a number of the impacts of malaria. The direct individual economic impacts of the disease include the value of lives lost, the value of time lost to sickness, and the expenditures on medical care, treatment, and prevention. Direct social costs include government expenditures on malaria control and prevention. The indirect costs may be greater still. These include changes in human settlement and labor patterns induced by the disease (e.g., changes in the locations where people live or farm). Indirect costs also include the consequences of the disease on fertility, demography, and human capital investments; on trade patterns and investment; and potentially on managerial quality and technology adoption. (For example, skilled managers may prefer not to work in malarial regions, resulting in reduced productivity levels.)[10]

Impacts of malaria on fertility and human capital decisions are difficult to identify, since all are related to income levels. The same is true of malaria's impacts on trade and investment and many other variables that are correlated with income levels.

In spite of the difficulties involved, two widely publicized papers have argued that malaria appears to slow economic growth in poor countries. Both papers use cross-country regression techniques and attempt to use instruments or controls to address the obvious causality problems. McCarthy, Wolf, and Wu (1999) find that malaria prevalence is negatively related to growth of per capita income. In turn, they find that malaria morbidity is linked to climatic differences across countries. The magnitude of

[9]UNICEF reports that the use of such bednets can reduce child mortality from malaria by 20 percent (http://www.unicef.org/health/index_malaria.html, accessed June 10, 2005).

[10]At an even more remote level, it might be possible to view human biological adaptations to malaria as part of the indirect cost. Thus, sickle cell traits, found in some individuals of African descent, are damaging and costly in their own right. Medical literature strongly suggests that the sickle cell trait confers some resistance to malaria and is thus an adaptive evolutionary response to the disease. Arguably, then, we could count the costs of sickle cell anemia as one of the indirect costs of malaria.

malaria's effect on growth is substantial: they find that Sub-Saharan African countries experience a reduction in income growth of 0.55 percent annually because of malaria. Using a relatively similar methodology, Gallup and Sachs (2001) find that countries with "intensive" malaria experience a reduction in per capita income growth of 1.3% annually. They suggest that, everything else being equal, a country experiencing intensive malaria would have its long-term level of income per capita reduced by one-third, compared with the same country in the absence of malaria.[11]

Based on this analysis, Sachs and other authors have suggested increasing current spending on malaria control by more than an order of magnitude. Global spending on malaria prevention and control is currently around $100-200 million annually.[12] But based in large part on his estimates of the economic impacts of the disease, Sachs (2005b) has estimated that $2-3 billion in annual spending would be needed to control the disease effectively in Africa alone. These larger sums are clearly within the capacity of the international community, but they would represent a substantial fraction of total aid disbursements by rich countries.[13] As a result, the increases would either require significant reallocation of existing aid portfolios or increases in the total quantities of foreign assistance given by rich countries.

To the extent that such increases in expenditure are justified by appealing to the likely impact on income levels and growth rates in malarial countries, it is useful to look further at the evidence for malaria's impact on income levels.

In the empirical literature, Acemoglu and Johnson (2007) offer a far more skeptical view of the growth effects of disease, based on an instrumental variables approach. They are joined in this skepticism by Weil (2007) and Cutler, Fung, Kremer and Singhal (2007). Some other authors (Bleakley 2009, Lucas 2005), however, find evidence that malaria eradication campaigns resulted – after very long time lags – in quantitatively significant impacts on health, fertility, and income. Some difficulties with the empirical literature include the paucity of reliable data and the inherent difficulty of identification.

To provide a different perspective on the issue, we find it useful to present a formal

[11] These numbers passed from academic research into policy; in the Abuja Declaration of 2000 40 African heads of state and governments signed on to a major commitment to fight malaria, citing these numbers as one major justification.

[12] Sachs and Malaney (2002) use the lower estimate.

[13] In 2002, OECD countries gave $58.3 billion in foreign assistance of all kinds.

model of malaria in a dynamic setting. Our paper is somewhat related to work by Chakraborty, Papageorgiou, and Pérez Sebastiàn (2007) that looks at an overlapping generation economy with malaria.[14] But because their analysis is based on a two-period model, quantitative results on the magnitude of disease impacts are hard to interpret.

Our paper brings an explicit dynamic general equilibrium framework to the question of malaria's impacts. We incorporate an epidemiological model of disease (following Gersovitz and Hammer 2004, 2005 or Philipson 2000), with a standard general equilibrium framework. Using a calibrated version of the model, we examine the impact of malaria on steady-state economic outcomes in the absence of prevention and control measures. We also model the impact of costly prevention measures, including measures that are less than fully effective.

3 The Model

We present here a heterogeneous-agent model that allows us to consider the link between income and disease prevalence. The model allows for endogenous determination of malaria infection rates, along with production levels and prices. Individuals are rational and forward-looking; however, they face idiosyncratic shocks (including disease shocks) that they are unable to contract away due to the incompleteness of credit and insurance markets. Essentially, this is a model in the spirit of Aiyagari (1994) and Huggett (1996), with some epidemiological features embedded. The epidemiological aspects of the model are similar to those presented by Philipson (2000), and we borrow from his analysis of "rational epidemics."

3.1 Model environment

The model environment has many individuals, born identical. New individuals are born each period. Some individuals die in each period, with mortality rates dependent on infection rates. Individuals are exposed to the disease in each period; some fall sick. The probability that an individual will become sick is positively related to the

[14] An earlier version of this paper is circulated as Papageorgiou, Chakraborty, and Pérez Sebastiàn (2005); this covers similar ground using a closely related model, but with an endowment economy that offers little insight into the interaction between income and disease.

fraction of individuals in the population who are already sick, so infection rates are endogenous to the model. Sick individuals face heightened probabilities of death and lower labor productivity. We model fertility rates as consistent with a constant population level. This implies that fertility rates will be higher when malaria is present, consistent with data..

Individuals are born healthy. They have zero initial asset holdings, but they accumulate assets through their lives. Assets can be rented to a representative firm in a perfectly competitive market for current-period production. However, there is no credit market, nor is there any insurance market. Therefore, individuals in the economy will use precautionary savings to protect themselves from idiosyncratic shocks. Assets vanish when people die.[15]

Note that some characteristics of this asset make it similar to human capital: people are born with no positive endowment; they cannot hold negative amounts; and their holdings disappear upon death. It is also the only savings technology in a rudimentary economy. In other respects, however, the asset is perhaps more like physical capital: it is measured in the same units as output and consumption, and thus it can be used to smooth consumption or to make "lumpy' purchases. In these respects, the asset is more analogous to physical capital than to human capital.

As in Philipson (2000), individuals may, at any point during their lives, make a lumpy purchase of a preventive good that will confer future protection from malaria. This lifetime prophylaxis requires a one-time expenditure of q units of consumption good. We think of this as the present value of a lifetime expenditure stream on bednets, drugs, and other preventive goods. Alternatively, if an effective vaccine were to become available for malaria, we could model this as the cost of the vaccine. Note that this is an indivisible purchase, and we initially model it as being totally and perfectly effective. In other words, once an individual has purchased the prophylaxis, he or she does not subsequently contract malaria, and there is no need for future spending. Subsequently, we will relax the assumption of perfect efficacy. In fact, our quantitative results, reported below, show that that when the preventive goods offer imperfect protection, there are large quantitative impacts on uptake, infection rates,

[15]This assumption effectively serves as a type of depreciation in the economy. We could equally well allow for assets to be redistributed to the new generation. The qualitative results of the model would not change significantly.

and economic outcomes. By assuming that q is lumpy, and by making it impossible for people to borrow, we bias our results towards making it difficult for people to afford the preventive good. In other words, we are biasing our results towards increasing the prevalence and economic impact of the disease.

3.2 Preferences and endowments

Preferences for any household i are given by the period utility function:

$$u\left(c_{it};\, s_{it}\right) = \frac{s_{it}\left[\gamma\left(c_{it} - \bar{c}\right)\right]^{1-\rho}}{1-\rho}$$

with lifetime utility given by: $\sum_{t=0}^{\infty} \beta^t u\left(c_{it};\, s_{it}\right)$, where s_{it} reflects a utility cost of being sick, such that $s_{it} \in \{\bar{s}, 1\}$, $0 \leq \bar{s} \leq 1$. A value of $s_{it} = 1$ corresponds to health, and a value of $s_{it} = \bar{s}$ corresponds to sickness. The parameter γ is a scalar that will determine the utility level of subsistence, and it will be calibrated below to give a plausible value for the "value of life" for people in the model economy, following Hall and Jones (2007).

Given their health status, households care only about consumption. They also face a subsistence consumption requirement, $\bar{c}$. This may be important in determining the affordability of disease prevention measures for different households.

Individuals are endowed with one unit of labor time in each period, which they supply inelastically to the labor market. Their effective labor units depend on health status, s_{it}, and π_{it}, which is an indicator of labor efficiency. This efficiency parameter is subject to idiosyncratic shocks and evolves according to a Markov process. Healthy individuals supply one raw unit of labor; if they are sick, however, their raw labor supply is reduced to $\bar{h}$. Effective labor units are determined by the raw labor supply and the idiosyncratic shock, so that:

$$h_{it} = \begin{cases} \pi_{it}\bar{h}, & \text{if } s_{it} = \bar{s} \\ \pi_{it}, & \text{if } s_{it} = 1 \end{cases}$$

Individuals have the capacity to influence their health status through the decision of whether or not to purchase prophylaxis against malaria. We define q to be a basket

of consumption goods necessary to achieve permanent disease protection (described in more detail below), and p_t is individual i's decision to purchase q. This choice is a binary choice, such that $p_{it} \in \{0, 1\}$.

Given this setup, the individual's period budget constraint is given by:

$$c_{it} + k_{i,t+1} + p_{it}q \leq w_t h_{it} \pi_{it} + r_t k_{it}$$

where $k_{it} > 0$ denotes accumulated assets, r_t is the return to capital, and w_t is the wage.

3.3 Technology

The technology side of our model economy is characterized by an aggregate technology with constant returns to scale. Individual effective labor units aggregate to $L_t = \sum_i h_{it}\pi_{it}$, and individual asset holdings aggregate to the physical capital stock $K_t = \sum_i k_{it}$. These are used to produce output Y_t according to the Cobb-Douglas production function:

$$Y_t = K_t^{\alpha} L_t^{1-\alpha}.$$

Factor prices then correspond to the marginal products of the factors. Thus, we assume that there is a perfect rental market for factors in this economy, with only spot markets available. Firms earn zero profits, and since there are no fixed costs, we can treat the economy as having a single cost-minimizing aggregate firm which rents capital and labor from the population and earns zero profits in equilibrium.

3.4 Population dynamics

In such an environment, population dynamics become important. We need to specify birth and mortality rates, which are differentiated across populations of sick and healthy people. We also need to model the risk of infection. Let d_h and d_s be the death rates of healthy and sick people, respectively. Let their fertility rate be f.

Defining N as the total population, we denote S as the proportion of sick people:

$$S = \frac{1}{N}\sum_i S_i, \text{ where } S_i = \begin{cases} 1, & \text{if } s_i = \bar{s} \\ 0, & \text{otherwise} \end{cases}.$$

Trivially, the proportion of healthy people in the economy can be written as $H = 1-S$. Let V be the proportion of people who have purchased prophylaxis. This is effectively a stock variable. In each period, there are also people purchasing prophylaxis; this fraction is given by:

$$P = \frac{1}{N}\sum_i p_i.$$

This group can in turn be divided into those who purchase when healthy and those who purchase when already sick. The healthy purchasers are given by:

$$P_h = \frac{1}{H}\sum_{i\in H} p_i,$$

while those purchasing q when already sick are given by:

$$P_s = \frac{1}{S}\sum_{i\in S} p_i.$$

Define the indicator variable $v_i$such that it takes a value of unity for individuals who have ever purchased protection and zero for all others. Then the fraction of individuals who are protected from disease is given by:

$$V = \frac{1}{N}\sum_i v_i.$$

Note that these individuals may be sick or healthy at the time when they purchase protection. Thus, we have $V_s = \frac{1}{N}\sum_i v_i S_i$ and $V_h = \frac{1}{N}\sum_i v_i\,(1 - S_i)$. In equilibrium, people who are sick will not choose to purchase protection, since it will not cure them of the disease. (We could model this differently, without any substantive change in the results.)

3.5 Laws of motion

Armed with this notation, we can write the laws of motion for different groups in the economy as follows:

For population, the net increment to population comes from deaths of sick and healthy people and from the fertility of sick and healthy people. Note that we do not treat men and women separately, nor do we model fertility rates as age-dependent,

so that all individuals in the model economy can bear children. Thus:

$$N' = N - d_s S - d_h H + fN.$$

The proportion S of sick people depends on births, deaths, and infection. Let I be the infection rate for healthy people who have not purchased prophylaxis. Then:

$$S' = \frac{N\left[S - d_s S + IH\left(1 - V\right)\left(1 - d_h\right)\right]}{N - d_s S - d_h H + fN}$$

The proportion of people who are protected from disease evolves according to the law of motion

$$V' = \frac{N\left[V - d_s V_s - d_h V_h + P_h H\left(1 - d_h\right) + P_s S\left(1 - d_s\right)\right]}{N - d_s S - d_h H + fN}.$$

We need still to characterize the infection rate I that applies for healthy people who have not purchased prophylaxis. Following Philipson (2000), we assume that the probability of contracting an infection depends on the proportion of people already infected and also on the inherent ecology of the disease. Thus, we make use of a formulation in which the infection rate itself evolves according:

$$i = Z\left(\frac{S}{N}\right)^{\mu}$$

where $\frac{S}{N}$ is the fraction of the population that is currently sick, Z is an index of malaria ecology, and μ is a parameter. This function has important properties. If either the population is fully healthy or the malaria ecology is zero, the next period's infection rate will be zero: this is a steady state. It is also the case that if both the infection rate and the ecology are at 1, this is another steady state. Note that our treatment of infection differs slightly from that of Philipson, whose "hazard rate" for infection combines both the natural rate of infection and the behavioral response. We define i here to be the probability that an unprotected individual will become infected in the next period; i.e., conditional on the individual not purchasing protection. Philipson's hazard rate, by contrast, gives an unconditional probability.

Finally, defining $C_t = \sum_i c_{it}$, and dropping time subscripts, the law of motion for

the aggregate capital stock is given by

$$K' = K + Y - C - PqN - d_s K_s - d_h K_h,$$

where K_s and K_h are respectively the aggregate capital held by the sick and the healthy. Note that the distribution of capital across individuals is non-degenerate. Indeed, good and bad luck with idiosyncratic shocks and health determine how much an individual accumulates. There is no borrowing or lending, nor is there insurance, so capital acts as a "rainy day fund" for individuals in the economy.

3.6 Equilibrium

We will define an equilibrium in this economy using a recursive approach. An equilibrium will consist of functions of the state variables for the economy and for the individuals:

- Functions for prices and wages;
- Functions for individual consumption, asset holdings, labor supply, and disease protection decisions;
- Distributions of health status and capital across individuals.
- Functions for the aggregate labor and aggregate capital employed in production, and the aggregate output produced;
- Laws of motions for each type's endogenous state

such that individuals of each type maximize utility subject to budget constraints, across states; the representative firm maximizes profits, subject to zero profits; factor markets and goods markets clear; the distributions of health status and capital are invariant; and the individual functions are consistent with the aggregate laws of motion for the economy.

Characterizing and solving for the equilibrium of this economy can be complicated. Note that disease dynamics imply that this economy will display multiple steady states. To see this, observe that with $S = 0$, there will be a steady state regardless of how many people purchase prophylaxis. In general, the existence of an interior

steady state ($0 < S < 1$) will depend on the cost of the protective goods, q, relative to the subsistence consumption requirement and the distribution of capital per person in the economy. With higher levels of capital, the economy can jump from one in which prophylaxis is generally unprofitable to one in which it is universal. Some poor economies, however, will never escape the high-disease trap. By contrast, other economies will start with sufficiently high levels of capital per worker that they will defeat the malaria burden.

The steady states here differ a little from those of standard Solow models. Like other models of this type, the steady state is determined as the point at which asset accumulation (initial asset endowments of the newly born plus savings from those who are alive) exactly offsets the loss of capital that occurs when individuals die. Our models display multiple steady states, because of the disease dynamics involved.

We view the multiplicity of steady states in the model as a substantively useful one for thinking about why some countries have been able to leave behind the problems of malaria, while other countries – even those with similar climate and geography – remain caught in a trap characterized by low productivity and high infection. For example, Singapore has effectively eliminated malaria infection, whereas Congo – a country that is reported to have a comparable malaria ecology – suffers from vastly higher rates of infection. Pakistan and Sri Lanka have roughly comparable malaria ecologies (Sachs et al. 2004), and income per capita in Sri Lanka is almost double the level in Pakistan (Heston et al. 2002), but Sri Lanka has a reported malaria prevalence rate that is 20 times that of Pakistan (Asian Development Bank 2005).

In our model, multiplicity allows for countries at similar income levels and with similar malaria ecologies to have different equilibrium levels of malaria prevalence, prevention, and other variables.

Modeling choices and implications

In a number of dimensions, we offer a model environment that will tend to exacerbate the impact of malaria on national incomes. Our model displays the following features that tend to maximize the impact of the disease:

- Once infected, people face reduced labor productivity and increased mortality probabilities for the remainder of their lives. (In reality, most episodes of the

disease last only one to two weeks, with few apparent aftereffects.)

- Protection measures must be paid for, up front, in lump-sum fashion; people cannot borrow to purchase these preventive goods, nor can they inherit money from their parents. (In reality, goods like bednets and drugs can be purchased as needed; here we assume that people must buy a lifetime supply, paying with their own savings.)

- People are born as working-age adults, so that all deaths affect productive workers, rather than dependent children.

- Fertility rates are assumed to match mortality rates, so that a country with malaria has proportionately more "young" people who have accumulated few assets. While this approach endogenizes fertility only in a mechanical sense, it does correspond to the observation that countries with high malaria burdens have large fractions of young individuals in the population, with correspondingly low levels of savings.

- By holding population fixed and abstracting from fixed factors, such as land, we avoid an obvious mechanism through which high mortality rates might actually improve income per capita of survivors.

All of these mechanisms tend to exacerbate the economic losses created by the disease. In some sense, then, the model offers malaria its "best chance" (or at least a very liberal opportunity) to have a negative impact on income per capita.

4 Calibration

We are interested in a set of quantitative experiments in which we assess the effects on aggregate output of various exogenous changes that will affect both malaria prevalence and economic variables of interest. To carry out these experiments, we need to select values for the parameters of the model. A number of the parameters we take from the literature, and others we choose to match observations for a stylized poor malarial country. For all the important parameters of the model, we perform robustness checks, as described below.

The parameters for preferences we take to be standard. The discount factor β we set to 0.95, assuming annual frequency, and we set the risk aversion parameter $\rho = 1$. The disutility of sickness is measured by the parameter $\bar{s}$, which we set equal to 1.0 in the benchmark economy, implying no disutility. We also report robustness checks using a value of 0.9, which is consistent with estimates of "disability weights" such as those reported by Murray and Lopez (1996). The change is not quantitatively important in our model.

Since malaria increases the probability of death in our model, we also need to consider the value that people associate with living as opposed to dying. For this we draw on estimates from the U.S. that estimate the statistical value of a life at approximately $4 million to $9 million (Viscusi and Aldy 2003). Taking $7 million as a reasonable middle number, we compute that this is approximately 11.3 times lifetime consumption in the U.S. As a result, we set the subjective value placed on living at 11.3 times annual consumption in the benchmark economy, which pins down a value of $\gamma = 11.3$. This number is also subjected to some robustness checks, which are reported below.

We use a value of 0.9 for the labor efficiency units of a person infected with malaria. This reflects a number of micro studies in the literature and is broadly consistent with Bleakley's work (2003) looking at malaria in the U.S. South. Since we assume that people who become sick with malaria never fully recover, in the model, we are essentially assuming that someone who falls sick with malaria faces a 10 percent loss in effective labor units, for as long as they live. As noted above, this tends to overstate the impact of the disease on labor productivity, relative to the data.

The subsistence constraint is set to zero in the benchmark economy.

Individuals also face idiosyncratic shocks independent of the risk of contracting malaria. We need to specify both the transition matrix for shocks and the magnitude of the shocks. In the experiments reported here, the magnitude of the shocks is taken to be 0.224 (following Domeij and Heathcote 2004), while the transition matrix is set to:

$$\begin{bmatrix} .900 & .100 \\ .100 & .900 \end{bmatrix}$$

We use a capital share on the aggregate production technology of 0.36, in keeping

with standard practice in the literature.

This leaves fertility rates and death rates for healthy and sick people, plus the crucial parameters relating to the cost of preventive goods and the infection rate.

Death rates are taken to be 0.075 for sick people (i.e., those infected with malaria) and 0.015 for healthy people. It is difficult to know which observations in the data to use for calibrating these parameters, but we believe the results to be quite robust to the death rates. To simplify the analysis, we set the fertility rate such that population will be stable in equilibrium. In other words, we allow the fertility rate to adjust to offset the deaths of sick and healthy people.

The cost of prophylaxis is another critical parameter for the model. Chima *et al.* (2003) provide a good summary of the literature on the costs of prevention and treatment of malaria in Africa. These numbers are hard to interpret, because (a) the figures given are often averages that include people who did not purchase preventive goods; (b) the goods on which people are spending money are not in fact effective in prevention (e.g., mosquito coils); and (c) the expenditure on bednets, screens, and mosquito coils is only partly intended to reduce malaria incidence, while also serving the purpose of reducing the annoyance of mosquito bites. Nevertheless, some reasonable numbers come out: bednets cost between \$5 and \$10 per person and last perhaps five years under reasonable use. At an interest rate of 0.05, the present value of a lifetime stream of bednet purchases at \$5 is about \$20-\$25 per person, which assuming per capita income of about \$500 could be modeled as a one-time fixed cost of 4-5% of annual per capita income. At \$10 per bednet, obviously, the number rises to 8-10% of per capita income. (At an interest rate of 0.10, this falls back to 5%.) The estimates of eventual vaccination costs are not much different in NPV terms, with estimates of \$20-\$60.[16] Thus, it seems that realistic values for this cost might range from 0.05 to 0.10 of annual income.

Finally, we have the parameters Z and μ for the infection rate process. Using the malaria ecology index of Sachs et al. (2004), we re-scale to define the index on the interval [0,1] and then find a value of 0.7 for a "typical" malarial country. In the data, this corresponds roughly to the level prevalent in Cambodia, Mozambique, Guinea-Bissau, or Congo. For that matter, it is also the malaria ecology prevalent in Singapore, a country with essentially no malaria. Thus, the malaria ecology value that

[16] At present, of course, no effective vaccine is available.

we choose is consistent in the real world with both malarial countries and malaria-free countries.

The parameter μ gives the elasticity of next period's infection rate with respect to the malaria ecology. We can estimate this by regressing infection rates on malaria ecology. A value of 0.122 was obtained from this regression.

5 Experiments and Results

Using the calibrated model, we conducted a number of experiments that we report below. The first experiment considers an economy in which protection from malaria is not possible. The first such experiment is to ask simply how large an effect malaria can have in an economy where no protective measures are available; in other words, where there is no behavioral response that is effective in reducing the burden of malaria. Arguably, this is a useful framework for thinking about the impact of the disease in some of the most severely affected environments, where neither spraying nor chemoprophylaxis nor drug treatments are effectively able to reduce the proportion of people suffering from the disease.

Specifically, the experiment we conduct is to calibrate the model to a set of benchmark parameters and then to suppose that the cost of a preventive bundle of goods, represented by q, is prohibitive.

We compare its healthy and unhealthy steady states. The second experiment considers the same question for an economy in which malaria protection is available, though costly. We carry out this experiment for a large range of possible costs. Finally, we repeat the second experiment for a range of possible parameter values, to assess the robustness of our results.

5.1 Experiment 1: A Benchmark Economy

The first experiment that we consider is one in which we compare the benchmark economy in two steady states, one of which has everyone healthy and the other of which has essentially all people sick. A simple way to arrive at these steady states is to set the cost of the preventative good at a very high level, so that it is effectively unavailable. Both steady states are feasible, and initial conditions in the model

economy will determine which one pertains. An economy that begins poor and sick will tend to say poor and sick, while one that starts with better health will end up at a better steady state.

The comparison of these two steady states offers an insight into the maximum possible impact of the disease within the model economy. In effect, we are examining the case in which there is no behavioral response to malaria. This provides a kind of upper bound of the disease's impact, within the model.

Table 1 shows the results of this experiment. The impact of the disease in this case is large. The steady state with widespread malaria infection has an income per capita that is 43 percent lower than that in the healthy steady state. Per capita consumption is even lower, with a 49 percent reduction from the healthy steady state. The proximate cause of the reduction is that steady state asset holdings are only 25 percent of the value in the healthy steady state. This reflects the shorter average lifespans of people in the malarial steady state: they do not live as long as those in the healthy steady state, nor do they expect to live as long, and they are poorer while alive. As a result, they save at a lower rate and accumulate assets over a shorter period. Figure 1 shows the distributions of asset holdings for the healthy steady state and the malarial steady state, and the impact of the disease is evident. It is this effect, rather than the direct impact of the disease on effective labor units, that has the greatest impact.

One way to think about this experiment is to view it as the benefits to the model economy of escaping from its malarial steady state and moving to its healthy steady state. Are there substantial impacts on income, as Sachs seems to suggest? Can we identify important *ex post* differences between the two economies?

The answer here is that there is a large difference in steady-state incomes between the two economies. The difference is not sufficient to explain why malarial countries are poor and non-malarial countries are rich, but it is true that in the model economy, eradication of malaria would lead approximately to a doubling of per capita income. This is a large impact. Note that the macro impacts of the response are far greater than the micro impacts; although individuals lose only 10 percent of their labor productivity to the disease, the lower asset accumulation leads to an amplification mechanism through which the disease impacts are multiplied.

5.2 Experiment 2: Costly Prevention

In the second experiment, we ask again about the impact of a single economy moving between its "healthy" steady state and its "sick" steady state. In contrast to the first experiment, however, we assume that an effective preventive good is available, though costly. Table 2 reports the results for a cost of prevention approximately equal to 25 percent of steady-state annual income; this would be comparable to a one-time cost of prevention of around $90, in an economy in which annual income is about $1 per day.[17]

What is the quantitative impact of the disease when costly but effective prevention measures are available? Table 2 shows the results of the experiment, comparing outcomes across the low and high steady states for a model economy in which $q = 0.6$. In this experiment, the economy in its low steady state has essentially the entire population protected from malaria, even though the protection is quite costly. People are willing to spend a large fraction of income to avoid getting sick.

The disease is not entirely eradicated, however, and people cannot forego the costs of protection. Indeed, as long as some malaria is present, there are individuals who become infected as newborns before they are able to buy prophylaxis. Clearly this is rare, but it does imply that the steady state of this economy gives slightly lower welfare one in which the disease is actually eradicated, where no one needs to bear the cost of the preventive good. Asset holdings, production, and consumption are all slightly lower than in the steady state with no malaria.

This result – that malaria matters only little – is at first sight surprising, given that we have given the disease every chance to have a major impact. The cost of lifetime protection is substantial and must be paid in full up-front (i.e., there is no borrowing to finance prophylaxis); agents are born with no assets; and they must also hold capital for precautionary savings. The risks of infection are low, since others are generally healthy. Yet even so, individuals in this economy are still willing to pay for protection as soon as they can afford it, and they afford it rapidly.

This seems to cast doubt on the potential for the disease to cause large macro effects in reality. Why would people in endemic areas not behave like individuals in

[17] This is not far from the expected lifetime costs of a vaccine that has to be readministered at five-year intervals, or from the total expected lifetime cost of insecticide-treated bednets, as calculated by Johnson (2007).

the model? Even if the individual costs of the disease are modest, would it not pay for people to purchase bednets or screens or drugs to prevent or treat malaria for themselves or their children? Our model seems to indicate that the disease should have little macro impact where there are effective protective measures available, even if they are somewhat costly.

5.3 Experiment 3: Varying the Cost of Prevention

Consider now the effects of varying the cost of protection from malaria. How large does the cost need to be before it is viewed as effectively unaffordable or undesirable (as in Experiment 1)? Figure 2 graphs a discrete approximation of the relationship between protection costs and the steady-state levels of output, consumption, assets, and the proportion of people sick and protected. As the protection cost rises, steady-state output falls in a weakly monotonic fashion.

A crude rule of thumb is that, for values of q less than one year's average income, essentially everyone in the model economy purchases the preventive good and buys protection from the disease. For costs much higher than one year's average income, some people opt not to purchase protection. Typically, these are individuals who have accumulated little capital and have had bad draws of the persistent idiosyncratic shock. For costs greater than twice the steady-state average income, essentially no one buys protection, including the "lucky rich." As a result, the economy faces the full force of the disease.

Thus, for malaria to have a big impact on income per capita, it must be true either that (a) there is not a truly effective bundle of preventive goods or actions; or (b) people are not aware of the effectiveness of the preventive goods; or (c) the cost of the preventive goods or actions is very high – in excess of one year's annual income. The model suggests strongly that a moderate charge for bednets or spraying or drugs, if these prevention and control measures were truly effective, would not deter people from purchasing these goods, simply for their private benefits.

Panel (d) of Figure 2 shows the proportions of people sick and protected as a function of the cost of q. Clearly the proportion sick rises as q increases, while the proportion buying protection falls (consistent with the law of demand).

Are the results of Experiment 3 driven by specific parameter values? Figures 3

and 4 demonstrate the robustness of the basic results to changes in the impact of the disease on the full range of parameters. It is striking that the qualitative results easily survive reasonably large changes in the parameter values.

5.4 Experiment 4: Limited Efficacy

Experiments 2 and 3 seem to raise a puzzle. Why is it that in actual malarial economies, relatively few people seem to use the protection measures that are available? In most countries, bednet use is very low, and the private demand for indoor residual spraying is even lower. In the model economy, by contrast, people seem willing to pay for preventive measures even when they are expensive. Why do individuals in malarial countries not take greater advantage of the available preventive goods? One hypothesis might be lack of information; another might be limited availability of the necessary items. But in most malarial countries, the basic preventive goods are widely available, and they are well understood, since government and non-governmental programs have been promoting their use for many years and, in some cases, even giving them away.

In this experiment, we ask whether a possible explanation might arise from limited efficacy. Hitherto, we have assumed that an individual who buys the preventive bundle is fully protected for life. In reality, however, the available protective goods are far less than one hundred percent effective. For example, bednets (which are currently the form of prophylaxis preferred by many policy makers) cannot protect people around the clock.

Our model predicts that if the bundle of preventive goods is less than fully efficacious, there will be a dramatic reduction in the fraction of people purchasing protection. The intuition behind this result is simple enough; the only thing worse for people in the model than getting sick would be to get sick after buying the preventive good.

Figure 5 shows the rapid drop-off in the fraction of people purchasing protection, for the cases where the protective bundle is less than fully effective. It is striking that even a relatively modest loss of efficacy would have large impacts on the economy. For example, at the benchmark level of prevention cost – of q equal to approximately one-fourth of annual income – a reduction from 100 percent efficacy to 99 percent efficacy

would have a very small but measurable impact on the fraction of the population purchasing protection. At a higher prevention cost of q equal to two years' income, however, a reduction from 100 percent efficacy to 99 percent efficacy would induce a decline of ten percentage points in the fraction of people purchasing prevention. A decline to 95 percent efficacy would reduce by half the number of people purchasing protection.

The effects on steady-state output would also be large. Figure 6 shows how steady-state output would be affected by decreases in the efficacy of the preventive good, holding all else constant. Again, at the benchmark prevention cost of one-fourth of annual income, a decline from perfect efficacy to 99 percent efficacy would reduce steady-state output by about seven percent. With q equal to two years' income, a decline from perfect efficacy to 99 percent efficacy would lead to approximately a ten percent drop in steady-state output. With $q = 2$, a decline to 95 percent efficacy would reduce steady-state output by about one-third. In part, these decreases stem from the lower proportion of people seeking protection from the disease. But another part, especially visible at low or zero protection costs, comes from the risk of getting sick despite purchasing the protection. Even a slight imperfection in efficacy, by the power of compounding, ends up having a major impact.

Although the model is clearly stylized, the analysis of efficacy has potentially important implications for policy. Where malaria prevention and control methods are less than fully effective – and the measures available today most probably have lower rates of effectiveness than the numbers analyzed here – it is to be expected that take-up rates will be very low, given any significant costs. Low take-up rates would be rational in this case, rather than reflecting ignorance or lack of information.

6 Conclusions

These results point to several notable conclusions. First, it is entirely possible for an economy to arrive at a "malaria trap," in which sickness begets poverty and poverty makes disease prevention unaffordable. In the model economy, we can quantify the magnitude of this "malaria trap." It can reduce income per capita by about half. By point of comparison, Gallup and Sachs (2001) note that the 44 countries with intensive malaria burdens in 1995 had per capita income of $1,526, compared with

$8,268 for the 106 countries without intensive malaria burden. Our model suggests that the disease alone could account for at most one third of this income gap.

However, the results also suggest that the impact of malaria on per capita income and consumption falls rapidly if people have the ability to protect themselves through behavioral responses. In fact, we find that the private demand for malaria prophylaxis is extremely strong – all the more so if the disease imposes severe economic consequences on its victims.

This raises a puzzle. Simply put, we find that in our model, three common claims about malaria are not together consistent with the observation that the disease has large impacts on per capita income. These claims are:

1. That malaria is a disease that imposes large private costs – in terms of health and economic consequences – for its victims;

2. That malaria can be controlled effectively by interventions such as bednets and drugs (either preventive or curative); and

3. That these prevention methods are perceived to be effective and are not prohibitively costly to those people at greatest risk of contracting the disease.

In our model economy, people are willing to spend large fractions of their income on prevention. As a result, if these preventive measures are reasonably effective, we find it difficult to generate scenarios in which the disease has big impacts on income per capita (or on welfare). For large aggregate impacts, it must be the case either that the preventive measures are unaffordable or ineffective – or else that people do not perceive the private costs of malaria as particularly large.

We do not believe that our results are uniquely dependent on the structure of our model; in fact, we believe that we have formulated our model in such a way as to maximize the likelihood that malaria affects aggregate output. In fact, it is hard to conceive of a model in which rational and forward-looking individuals would fail to adopt prevention measures in the face of the three claims listed above. Thus, we are inclined to conclude that at least one of the three claims does not hold in reality.

We offer no empirical insight as to which of the three claims might be false. Some evidence suggests that the private economic costs of malaria are modest, at least to adults. (Many adults in endemic areas eventually develop partial resistance to

malaria and experience it as a relatively mild flu-like disease on an occasional basis). But we feel that this viewpoint dramatically underestimates the losses of utility that come from the sickness and death of infants.

A second possibility is that users are poorly informed about the effectiveness of the available prevention and control measures, failing to appreciate their value. This is a view expressed frequently by field workers involved in the distribution of bednets and occasionally of drugs. But again, we believe that the private demand for malaria prevention should be quite strong, given any plausible accounting for the private costs. If a bednet costs two days' wages and yields a substantial reduction in infant mortality, and if people are unwilling to pay this cost without subsidy, it raises puzzling questions about the implied valuation of life.

A third possibility is that the available methods of prevention and control are, in practice, less effective than their advocates argue (or at the very least, that people believe these methods to be somewhat ineffective). While we offer no specific evidence for this hypothesis, we do show that in the model economy, private demand for the preventive measures falls rapidly if their effectiveness is less than perfect. We find that in the model economy, the disease can have significant impacts if the available methods of prevention are less than fully efficacious.

We calculate that in a poor tropical economy with a severe malaria burden, there would be little increase in income per capita if a low-cost vaccine were introduced that protected people from malaria with an effectiveness rate of 60 percent. In the same economy, however, an equally low-cost vaccine with perfect effectiveness could almost double income per capita.

Much recent debate has focused on the cost at which bednets and malaria drugs are distributed to people in endemic areas. Our analysis suggests that the costs of prevention must be very large, relative to average annual income, before they significantly affect the take-up of the preventive good. However, demand is far more sensitive to efficacy. The perceived effectiveness of the preventive good at stopping malaria matters a great deal; if people believe that they will become sick regardless of whether or not they purchase or use the preventive good, the demand will drop very quickly.

A point worth noting here is that all of our analysis looks at private responses to malaria. But we know that decentralized outcomes in this economy are not in general

optimal, because there are important infection externalities operating. Individuals in our model economy do not weigh in their decisions the potential impact of their actions on others. In particular, they are likely to under-invest in prevention relative to the social planner's optimum. This suggests that there may be a role for the public sector to undertake campaigns of prevention and/or treatment, or to subsidize the costs of preventive goods. But given that this cost has a minor impact on the uptake of prevention, within a plausible range, price subsidies may not be a priority.

It is important to note that the ultimate justification for investments in malaria control and treatment is the welfare cost, rather than the reduction in steady-state income per capita. Even if we found that the impacts on steady-state income were small, there are many other reasons why we should care about malaria and the enormous and tragic harm that it does. For hundreds of thousands of families, malaria is killing their infants and children. In many other families, the disease interferes with daily life, including schooling. Whether or not these effects are important for national income, they matter deeply to the individuals and communities that are affected. We do not need to justify malaria control programs on the grounds that they will contribute to GDP or to GDP growth. This must be one ingredient of our thinking, but the moral imperative alone is surely sufficient to justify some efforts for prevention, control and treatment.

7 References

Acemoglu, Daron and Simon Johnson. 2007. Disease and development: The effect of life expectancy on economic growth. *Journal of Political Economy* 115(6): 925-85.

Aiyagari, S. Rao. 1994. Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics* 109(3): 659-84.

Ashraf, Quamrul H., Ashley Lester, and David N. Weil. 2009. When does improving health raise GDP? *NBER Macroeconomics Annual 2008*, eds. Daron Acemoglu, Kenneth Rogoff, and Michael Woodford. NBER Books: National Bureau of Economic Research, Inc.

Asian Development Bank. 2005. *Key Indicators 2005: Labor Markets in Asia: Promoting Full, Productive, and Decent Employment.* Manila: ADB.

Bleakley, Hoyt. 2003. "Disease and Development: Evidence from the American South." *Journal of the European Economic Association* 1.2-3: 376-86.

Bleakley, Hoyt. 2009. Malaria eradication in the Americas: A retrospective analysis of childhood exposure. Forthcoming, *American Economic Journal: Applied.*

Chakraborty, Shankha, Chris Papageorgiou, and Fidel Perez-Sebastian. 2007. Diseases and Development. Mimeo: University of Oregon Department of Economics.

Chima RI, C.A. Goodman, and A. Mills. 2003. The economic impact of malaria in Africa: a critical review of the evidence. *Health Policy* 2003(63): 17-36.

Cutler, David, Winnie Fung, Michael Kremer, and Monica Singhal. 2007. Mosquitoes: The long-term effects of malaria eradication in India. Mimeo: Harvard School of Public Health.

Domeij, David and Jonathan Heathcote. 2004. On the distributional effects of reducing capital taxes. *International Economic Review* 45(2): 523-54.

Gallup, John Luke and Jeffrey D. Sachs. 2001. The economic burden of malaria. *American Journal of Tropical Medical Hygiene* 64(1, Supplemental): 85-96.

Gersovitz, Mark and Jeffrey S. Hammer. 2004. The economical control of infectious diseases. *The Economic Journal* 114(492): 1-27.

Gersovitz, Mark and Hammer, Jeffrey S. 2005. Tax/subsidy policies toward vector-borne infectious diseases. *Journal of Public Economics* 89(4): 647-674.

Hall, Robert E. and Charles I. Jones. 2007. The value of life and the rise in health spending. *Quarterly Journal of Economics* 122(1): 39-72.

Hammer, Jeffrey S. 1993. "The economics of malaria control." *World Bank Research Observer* 8.1: 1-22.

Hamoudi, Amar and Jeffrey D. Sachs. 1999. The changing geographical distribution of malaria: A review. CID Working Paper number 2. Cambridge, Mass.: Kennedy School of Government, Harvard University.

Heston, Alan, Robert Summers and Bettina Aten. 2002. Penn World Table Version 6.1, Center for International Comparisons at the University of Pennsylvania (CICUP).

Holding, P. A. and R.W. Snow. 2001. Impact of *Plasmodium falciparum* malaria on performance and learning: review of the evidence. *American Journal of Tropical Medical Hygiene* 64(1,2)S: 68-75.

Huggett, Mark. 1996. Wealth distribution in life-cycle economies. *Journal of Monetary Economics* 38(3): 469-94.

Johnson, Kelsey. 2007. Evaluating the cost effectiveness of malaria control strategies: A comparison of insecticide-treated nets, long-lasting insecticide nets, indoor residual spraying, and vaccination. Manuscript, Yale University.

Lucas, Adrienne. 2005. Economic Effects of Malaria Eradication: Evidence from the Malarial Periphery. Manuscript. Department of Economics, Brown University.

McNeill, William H. 1976. *Plagues and People.* Garden City, NY: Anchor Press. Reprinted by Anchor Books in arrangement with Doubleday, 1998.

Murray, Christopher J. L. and Alan D. Lopez. 1996. Evidence-based health policy–lessons from the Global Burden of Disease Study. *Science* 274(5288): 740 – 743.

Papageorgiou, Chris, Shankha Chakraborty, and Fidel Perez-Sebastian. 2005. Diseases and Development. Departmental Working Papers 2005-12, Department of Economics, Louisiana State University.

Philipson, Tomas. 2000. Economic epidemiology and infectious diseases, in: A. J. Culyer & J. P. Newhouse (ed.), *Handbook of Health Economics*, chapter 33, pages 1761-1799, Elsevier.

Sachs, Jeffrey. 2005a. *The End of Poverty.* New York: The Penguin Press.

Sachs, Jeffrey. 2005b. Achieving the Millennium Development Goals – The case of malaria. *New England Journal of Medicine* 352(2, January 13): 115-117.

Sachs, Jeffrey, Anthony Kiszewski, Andrew Mellinger, Andrew Spielman, Pia Malaney, and Sonia Ehrlich Sachs. 2004. A global index of the stability of malaria transmission. *American Journal of Tropical Medicine and Hygiene* 70(5): 486-498.

Sachs, Jeffrey and Pia Malaney. 2002. The economic and social burden of malaria. *Nature* (415: 7 February): 680-85.

Spalding-Fecher, Randall, and Shomenthree Moodley. 2002. "Economic valuation of increased malaria due to climate change: A South African case study." *South African Journal of Economic and Management Sciences, N.S.* 5.2: 395-412.

Tishkoff, Sarah A., Robert Varkonyi, Nelie Cahinhinan, Salem Abbes, George Argyropoulos, Giovanni Destro-Bisol, Anthi Drousiotou, Bruce Dangerfield, Gerard Lefranc, Jacques Loiselet, Anna Piro, Mark Stoneking, Antonio Tagarelli, Giuseppe Tagarelli, Elias H. Touma, Scott M. Williams, and Andrew G. Clark. 2001. Haplotype diversity and linkage disequilibrium at human G6PD: Recent origin of alleles that confer malarial resistance. Originally published in *Science*

Express on 21 June 2001 (DOI: 10.1126/science.1061573). *Science* (293: 5529, 20 July) 455-462.

Viscusi, W Kip and Joseph E. Aldy. 2003. The value of a statistical life: a critical review of market estimates throughout the world. *Journal of Risk and Uncertainty* 27(1): 5-76.

Weil, David. 2007. Accounting for the effect of health on economic growth. *Quarterly Journal of Economics* 122(3): 1265-1306.

Table 1: Experiment 1 Results (Multiple steady states with prohibitively expensive protection).

$q = 1000$ $z = 0.7$	Low	High
Endogenously determined fertility rate	0.0690	0.0150
Proportion sick	0.9007	0.0000
Proportion protected from disease	0.0000	0.0000
Average assets	2.9596	12.0797
Average output	1.3913	2.4521
Average consumption	1.1565	2.2668

Table 2: Experiment 2 Results (Multiple steady states with feasible but costly disease protection).

$q = 0.6$ $z = 0.7$	Low	High
Endogenously determined fertility rate	0.0150	0.0150
Proportion sick	0.0006	0.0000
Proportion protected from disease	0.9770	0.0000
Average assets	12.0631	12.0797
Average output	2.4488	2.4521
Average consumption	2.2553	2.2668

Figure 1. Distributions of individual asset holdings in benchmark economy, steady states with and without malaria.

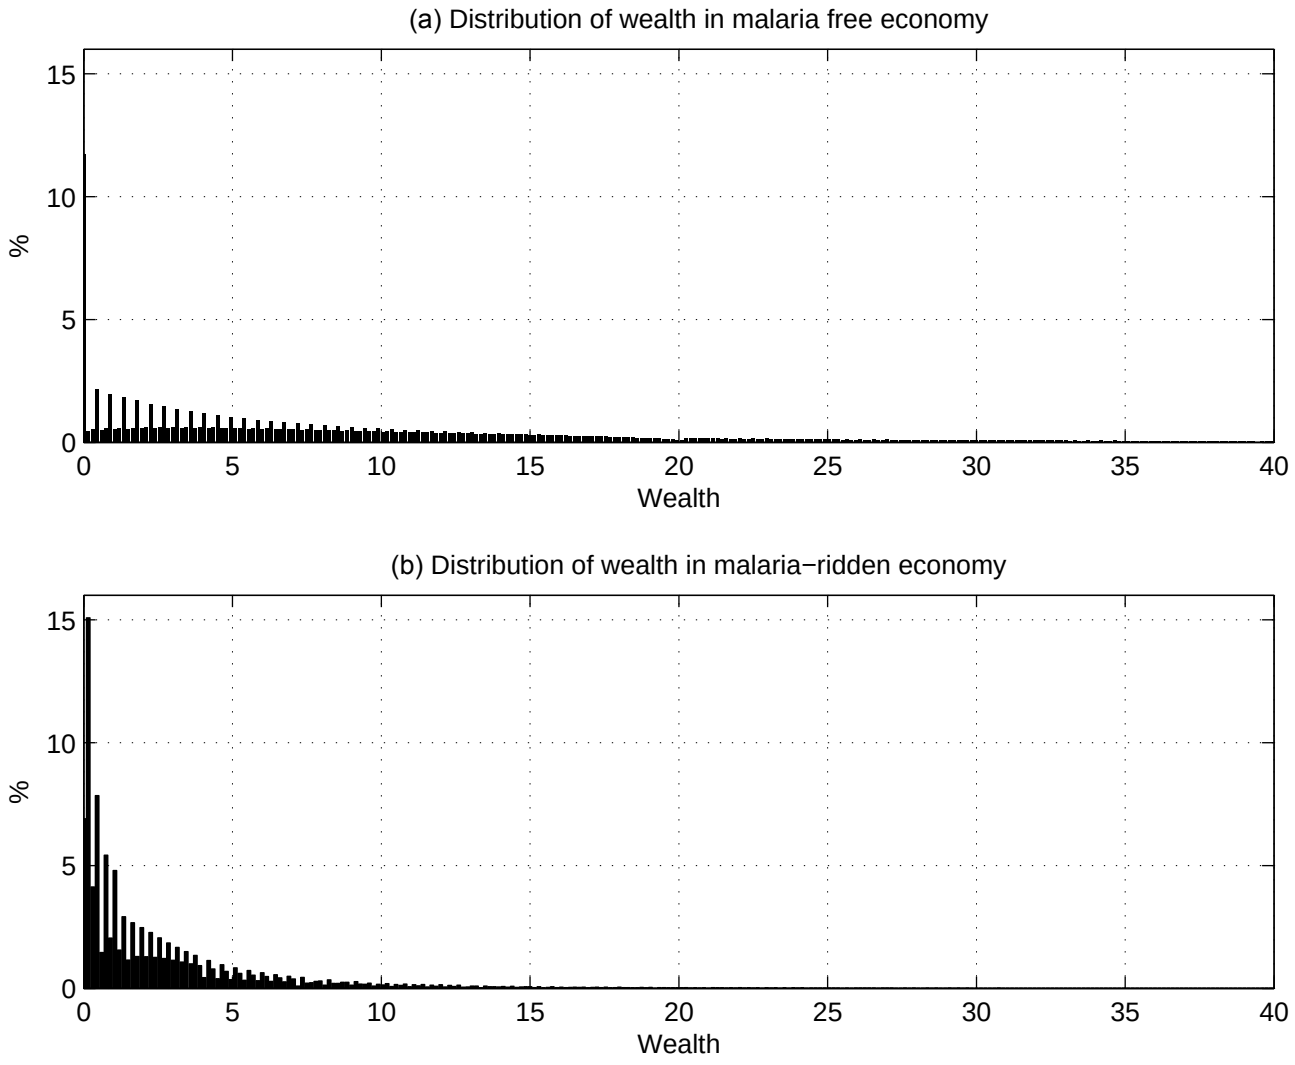

Figure 2: Benchmark economy, showing response of key variables to changes in the cost of protection from disease, in "sick" and "healthy" steady states.

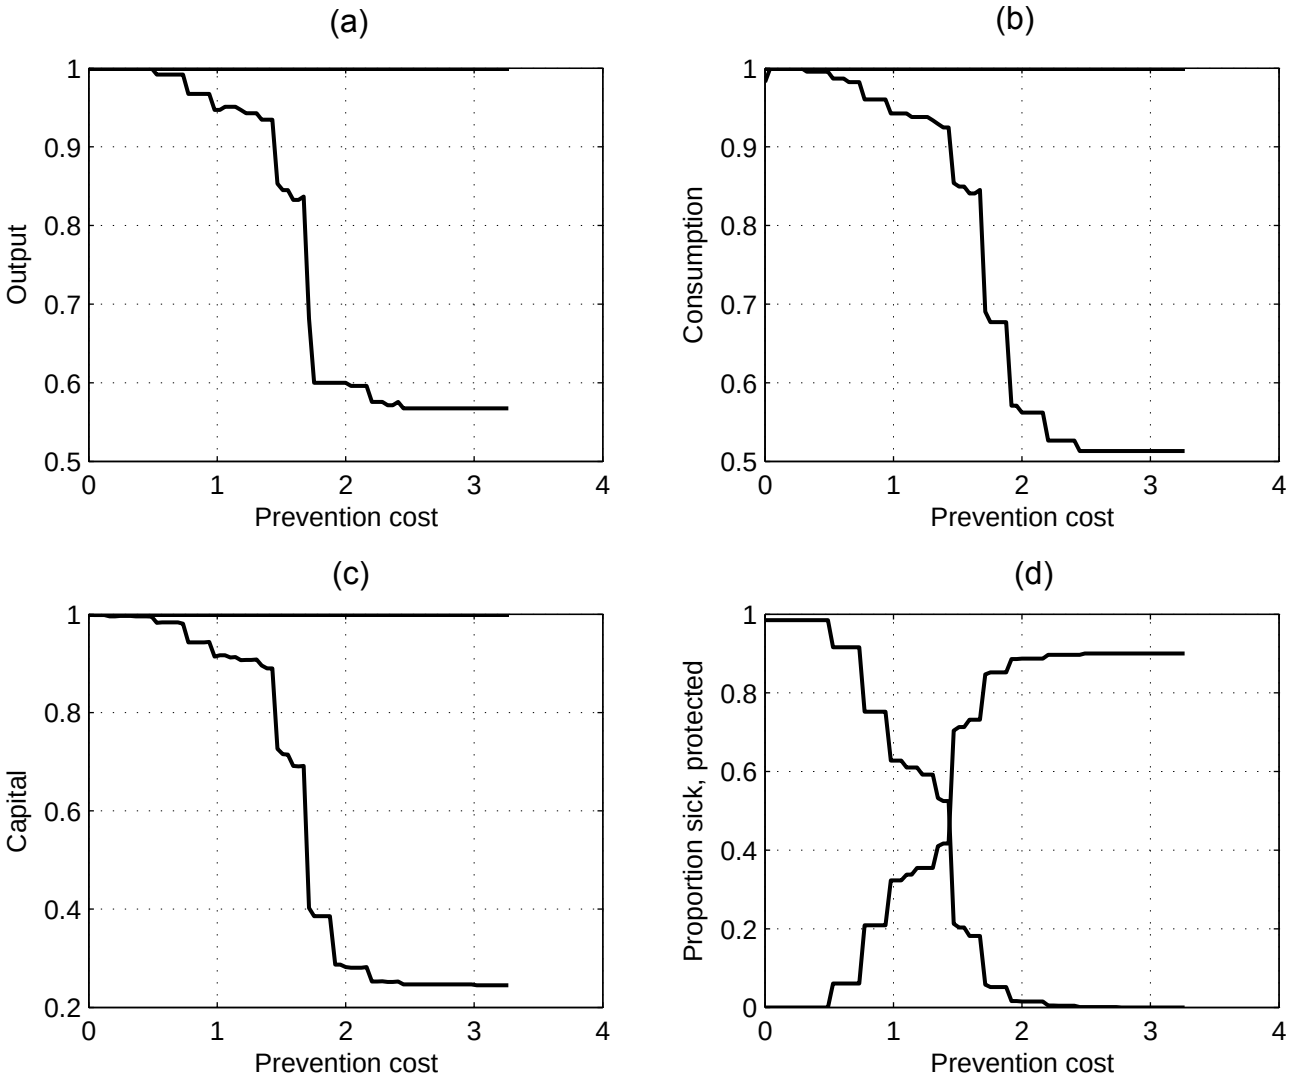

Prevention costs are measured relative to average income in the malaria-free economy.

Figure 3: Robustness checks -- sensitivity of the model to changes in parameter values.

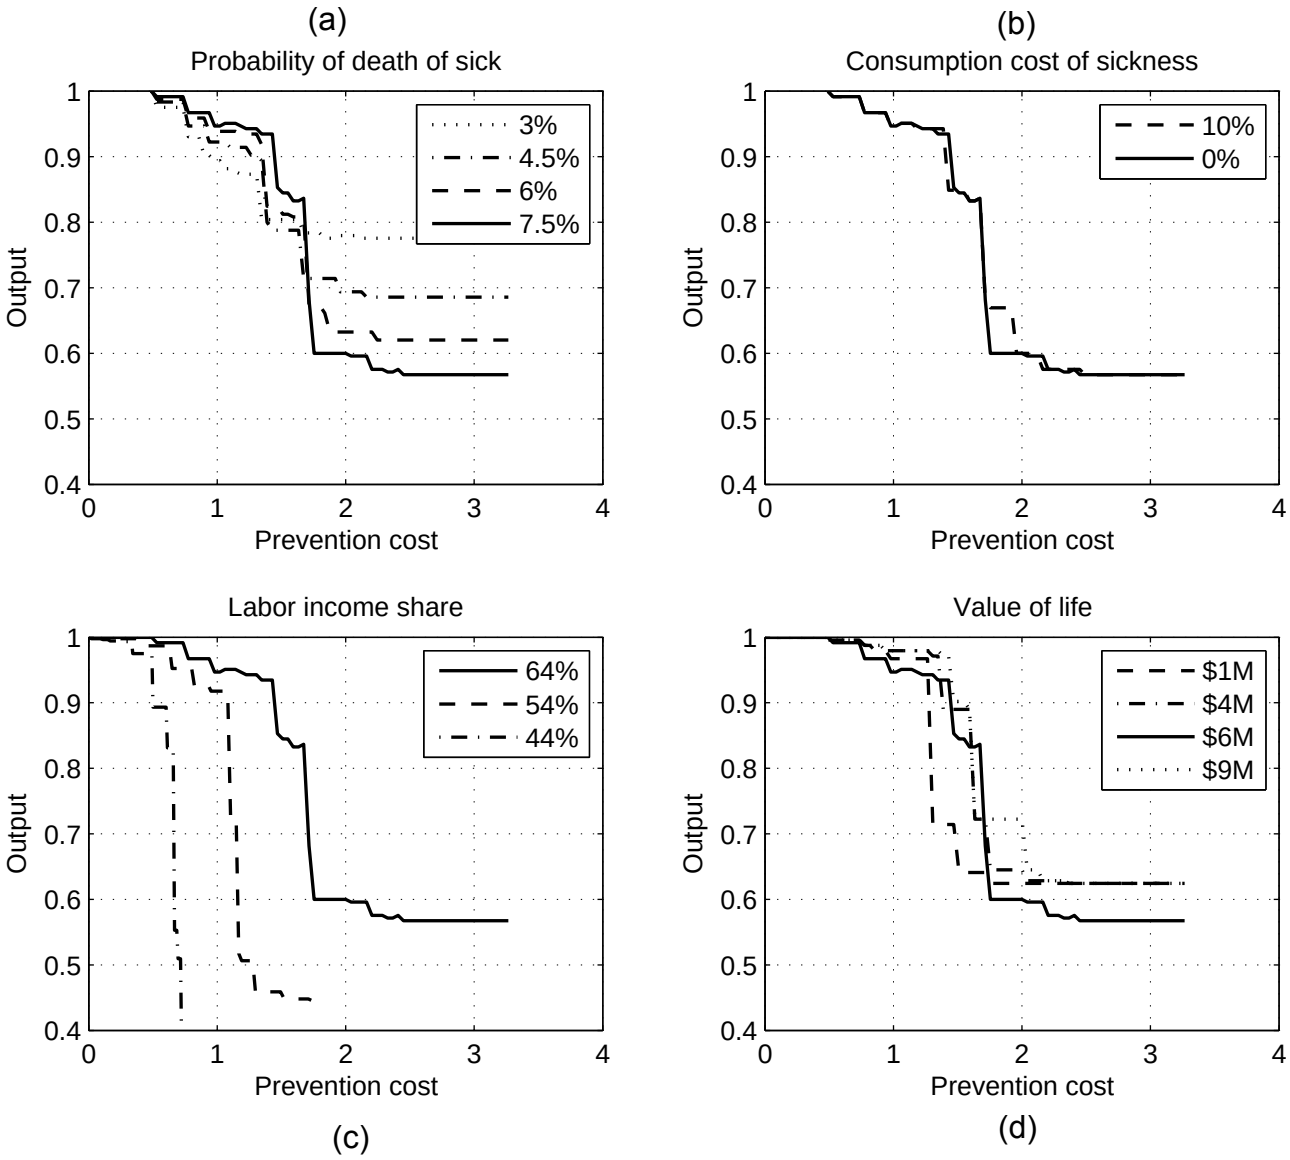

Prevention costs are measured relative to average income in the malaria-free economy.

Figure 4: Robustness checks -- sensitivity of the model to changes in key parameter values.

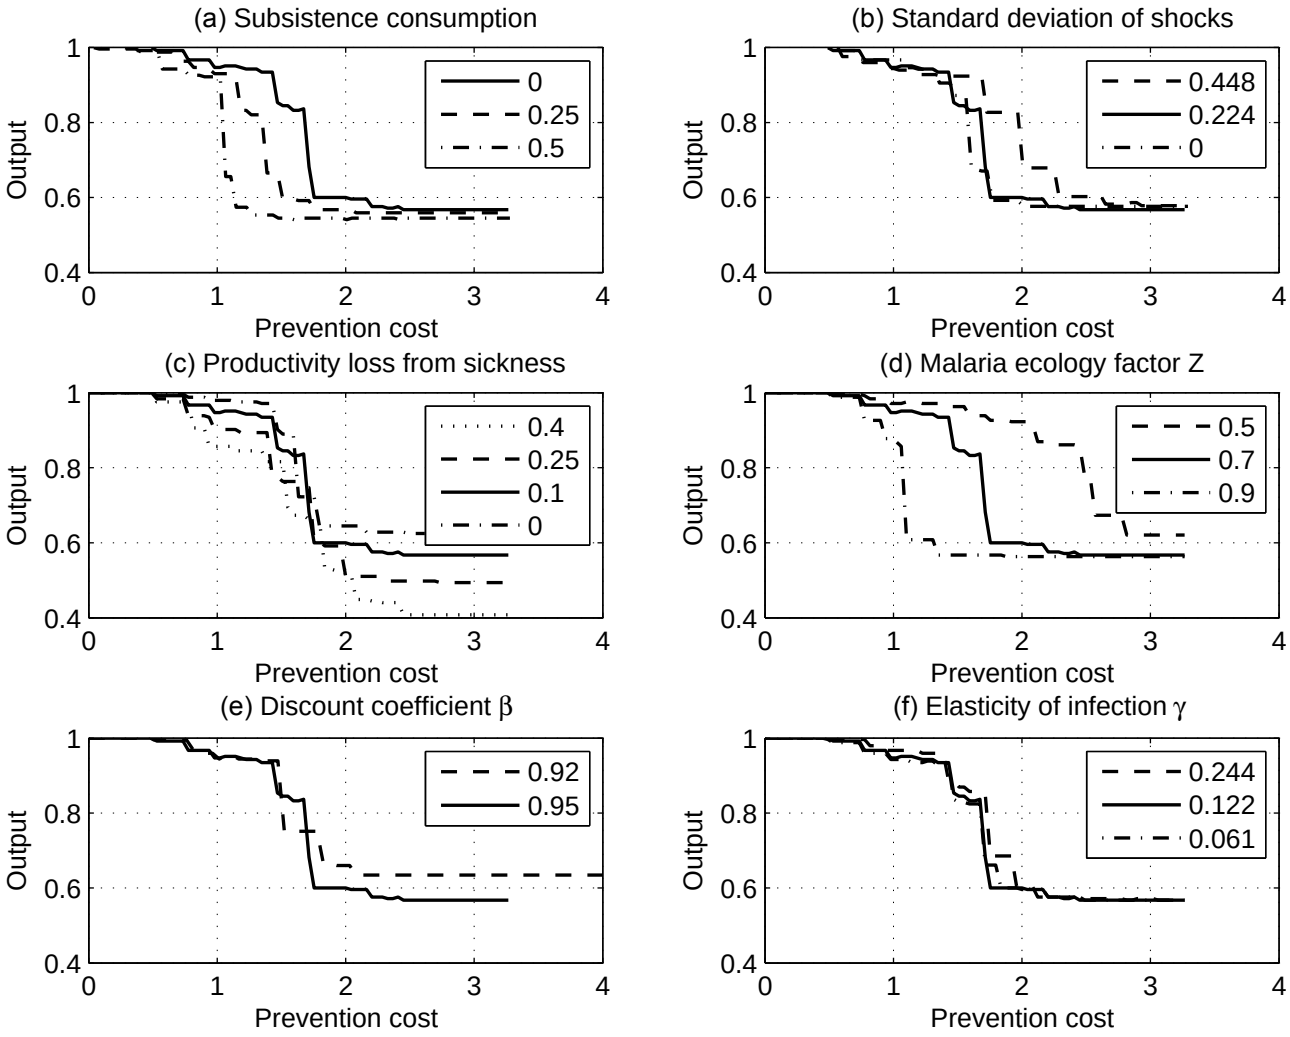

Prevention costs are measured relative to average income in the malaria-free economy.

Figure 5: Proportion of people who have purchased protection in steady-state, for economies differing in the degree of efficacy of the preventive good.

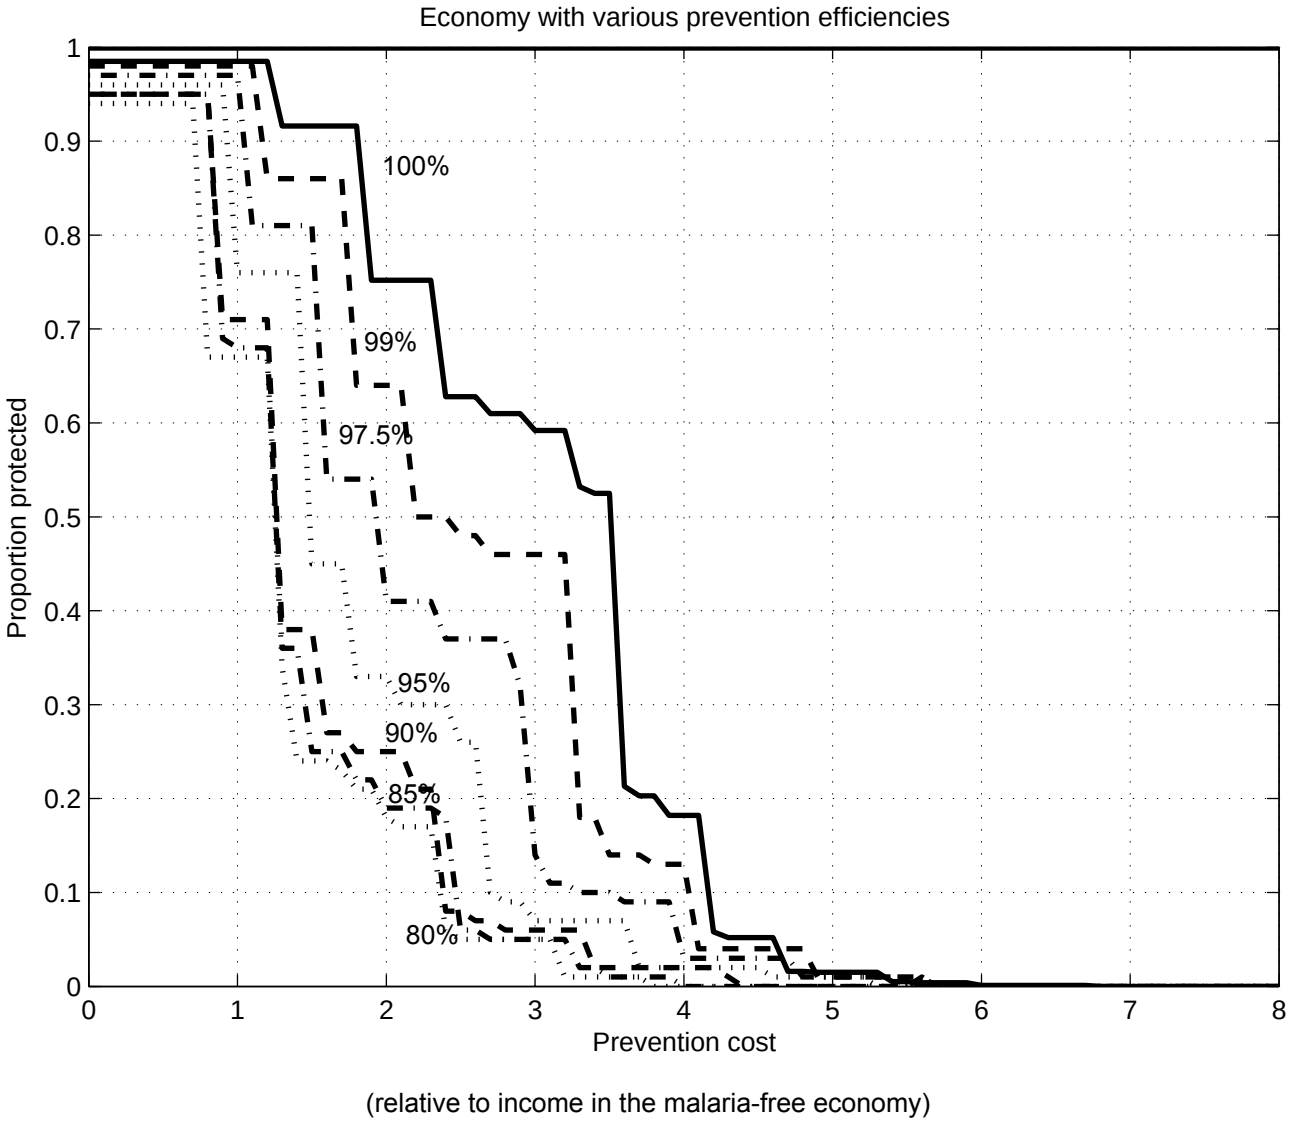

(relative to income in the malaria-free economy)

Figure 6: Output per person relative to benchmark economy,
for economies differing in the degree of efficacy of the preventive good.

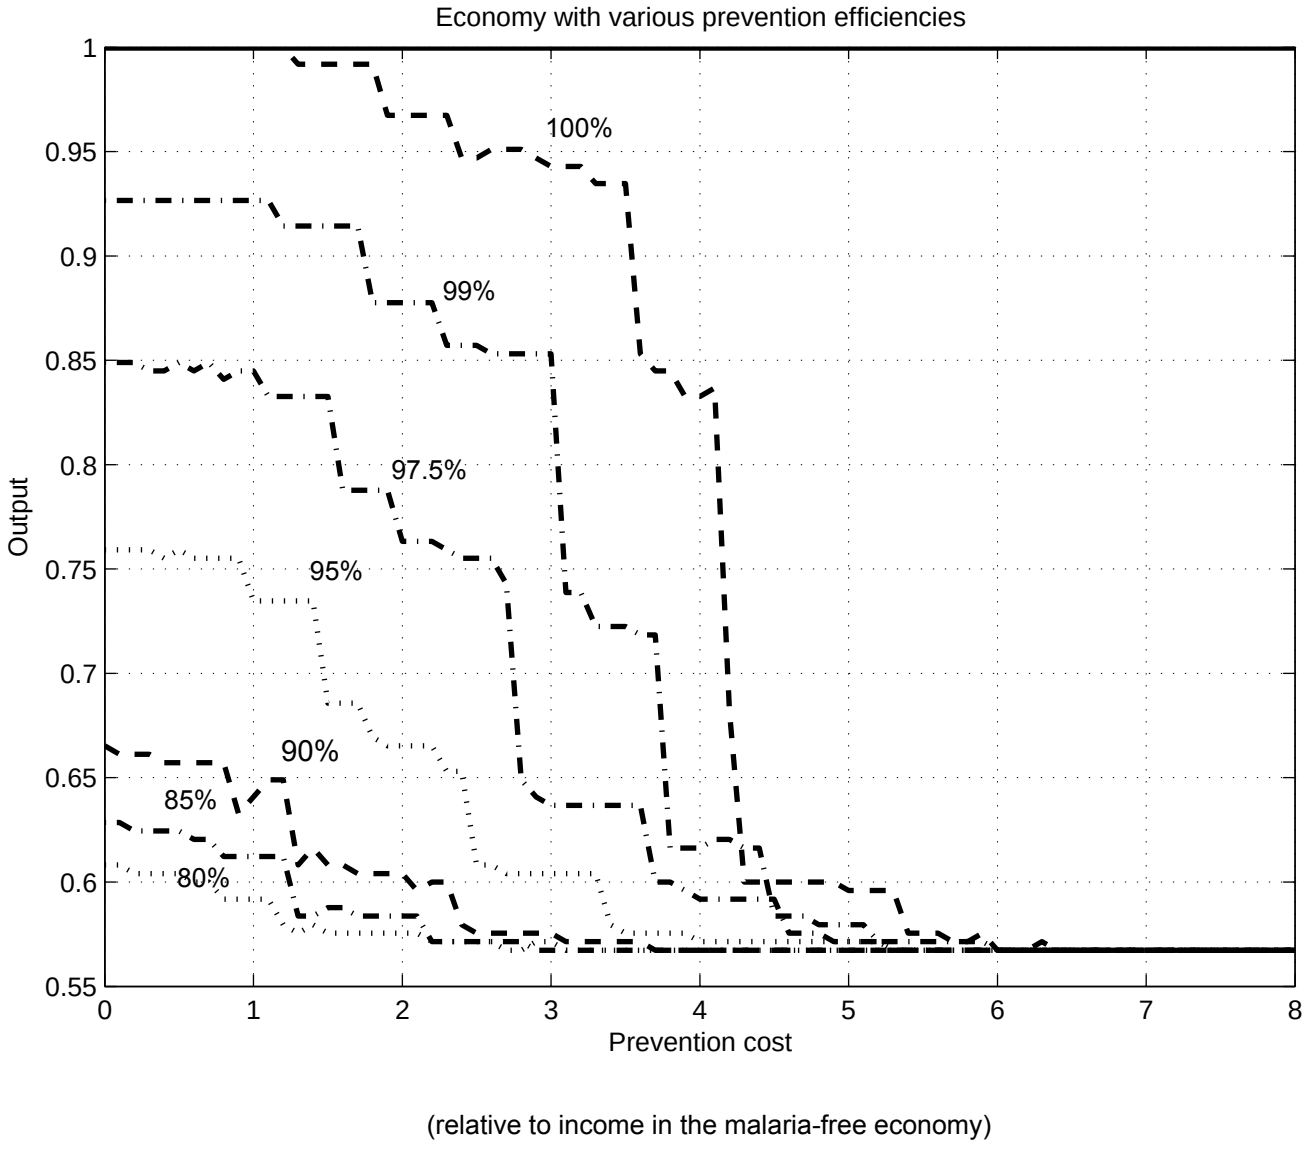

(relative to income in the malaria-free economy)