

Malaria Policy:

Alternative Prevention and Eradication Strategies in a Dynamic Model

Douglas Gollin
Williams College

Christian Zimmermann
University of Connecticut

First version: August 12, 2007

This version: September 3, 2007

ABSTRACT

As many as one million deaths annually are attributed to malaria, a parasitic disease transmitted by mosquitoes. The heaviest burdens of malaria are found in poor tropical countries, where medical systems are often ill-equipped to treat the sick, and where public health programs may lack the resources for effective and sustainable campaigns of prevention and control. In recent years, the international community has – perhaps belatedly – increased the funding for malaria programs. Some donors have invested heavily also in scientific research on malaria, such as vaccine and drug development. Others have advocated subsidies for inexpensive and available controls, such as distributing bed nets, or spraying insecticides and larvicides that would target mosquito populations. Relatively little quantitative analysis has guided these investments, however. This paper reports on an effort to compare alternative control and prevention strategies, using a dynamic model that incorporates some essential economic and epidemiological features. We compare steady-state equilibria produced by alternative treatment and control methods, and we use these to calculate the costs and benefits of different approaches to malaria control.

Authors' addresses: Douglas Gollin, Department of Economics, Williams College, Morey House, 121 Southworth St., Williamstown, MA 01267 (dgollin@williams.edu); Christian Zimmermann, Department of Economics, University of Connecticut, 341 Mansfield Road, Unit 1063, Storrs, CT 06269 (Christian.Zimmermann@uconn.edu).

1. Introduction

In recent years, the international community has embarked on an ambitious (and arguably belated) effort to reduce the prevalence and impact of malaria. The disease, which is carried by mosquitoes, is widely thought to cause as many as 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

After years of relative neglect, international funding has begun to pour into malaria research, prevention, treatment, and control efforts – spurred in part by claims that the economic impact of malaria is large. In 2005, the United States government announced a \$1.2 billion aid program aimed at reducing malaria deaths in Africa. The Gates Foundation has also allocated about \$800 million to malaria research and control programs, in addition to \$650 million that it has given to the Global Fund to Fight AIDS, Tuberculosis, and Malaria. The World Bank is pledging an additional \$500 million over three years for its Global Strategy and Booster Program.³

Perhaps surprisingly, however, these investments appear to have been guided by little or no systematic priority-setting exercises that would analyze the costs and benefits of different prevention, control, or eradication strategies. Given the magnitude of these sums being spent, it would seem worthwhile to ask how best the resources might be used to prevent or control malaria, or to alleviate the disease's effects.

For starters, how much would it cost to protect large numbers of people from malaria? Even the least expensive measures for prevention and treatment – such as drugs and bednets – might cost several dollars per person per year. It is unlikely that these methods would be sufficiently comprehensive to halt the cycle of transmission and eradicate the disease, so these expenditures would be required on a continuing basis. With as much as 40 percent of the world's population living in areas where malaria is endemic, a simple – and naïve – calculation suggests that it could cost \$5-10 billion annually to provide full funding of these

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.

³ It would be wise, however, to cast a skeptical eye towards these promises of high levels of aid for malaria. In a revealing paper, Narasimhan and Attaran (2003) show that donor organizations' claimed support for malaria control efforts may overstate the actual levels of expenditure by an order of magnitude. Nevertheless, the recent upswing in malaria control expenditure appears to be real.

measures for every affected individual. Compared to total official development assistance (ODA) of \$53.1 billion in 2000, this is a large number.

But would expenditures of this magnitude actually be required? This figure overestimates the public costs of malaria prevention and control in at least two ways. First, many individuals may display high willingness-to-pay for prevention and treatment, due to the high private returns. These people would be willing to pay part or all of the cost of effective malaria controls/treatments. Second, in many countries and regions, it is unnecessary to achieve 100 percent coverage. Significantly smaller coverage rates may be sufficient to disrupt the transmission mechanism for the disease.

This paper attempts to assess both the private and public costs of reducing or eliminating malaria in a model economy that is designed to mimic some essential features of a developing economy. Implicitly, we seek to ascertain which of a set of potential control strategies will be most cost-effective. Will it be more cost effective to offer a partial subsidy on prevention and treatment measures, or to bear the full costs of these measures for a small group of people, to ensure that transmission mechanisms are disrupted, or to pursue a policy that mixes these two strategies?

To address these questions, we use a calibrated dynamic general equilibrium model that extends our previous work (Gollin and Zimmermann 2007) on the aggregate impacts of malaria. This paper in turn follows Chakraborty et al. (2007), Gersovitz and Hammer (2004, 2005), Hammer (1993), and Philipson (2000) in applying an explicit epidemiological model to the economic analysis of disease.

In our model economy, individuals who are infected with malaria face declines in labor productivity and increases in mortality risk, with corresponding losses in expected lifetime income and well-being. People have access to preventive measures – comparable to bednets and spraying for mosquitoes, or to a hypothetical vaccine – but these are costly and may be imperfectly effective. Individuals would choose to purchase preventive measures whenever the expected private benefits of prevention exceed the costs. However, some individuals who would like to buy protection are unable to do so, due to credit constraints. (In the model, individuals cannot borrow to purchase malaria prophylaxis.) Moreover, when individuals are deciding whether or not to purchase prophylaxis, they do not take into account the impact of their preventive behavior on others in the economy. For both reasons, there tends to be an under-provision of malaria prevention activities in the model economy, relative to the social optimum. Thus, a role for policy emerges.

Our paper does not solve for the optimal policy, in a Ramsey sense, but the model allows us to compare the costs and impacts of a finite set of alternative interventions in the market for the preventive good. For each policy, the model identifies a corresponding steady-state level of output and prices, along with morbidity and mortality. There are no direct analytic solutions, but we can calculate a full set of computational results that describe the model economy under various policies.

The outline of our paper is as follows. Section 2 provides some further background on malaria and possible control measures. Section 3 introduces the model, which (as noted above) draws heavily on our previous paper (Gollin and Zimmermann 2007). Section 4 describes the calibration of the model, and Section 5 reports the results of a set of quantitative experiments conducted in this model economy. Finally, Section 6 concludes.

2. Background

Malaria is an ancient disease that may in fact have co-evolved with early humans over a period of hundreds of thousands of years (Carter and Mendis 2002). Certainly it has been recognized as a disease for at least 2,500 years, with its symptoms being well described in some of the earliest medical writings known, from India, Greece and China (Carter and Mendis 2002). For many years, however, the transmission mechanism of malaria was poorly understood. thought to occur through bad air (“*mal aria*”), and only in the 19th century did the French doctor Laveran identify the parasitic nature of the disease. Ronald Ross, an officer in the colonial Indian Medical Service, identified mosquitoes as the vector of transmission in 1897.

Several traditional plant-based remedies were used to treat malaria in different parts of the world. In China, sweet wormwood (“qinghao”) was recommended as a treatment for malaria as early as the third century of the modern era. The bark of the tree *Quinina calisaya* (“cinchona”) plant was used as a traditional treatment in the Andean region. The active compound, an alkalid, was later isolated and called quinine. Drugs based on sweet wormwood and quinine remain the first-choice pharmaceutical treatments for malaria today.⁴

The identification of the anopheles mosquito as the vector for malaria led to a different set of approaches to prevention and eradication of the disease. As early as 1900, Patrick Manson attempted to show that sleeping under mosquito nets

⁴ The history of malaria and its treatments can be found from many sources; a useful concise reference is the web site of the U.S. Centers for Disease Control, which maintains a web site on the history of malaria: <http://www.cdc.gov/malaria/history/index.htm> (August 2007).

would provide protection from malaria. Ross himself initiated efforts to reduce mosquito breeding sites through draining pools and puddles, and by the early 1900s, the British authorities in India were using larvicides to limit mosquito populations. Spraying of insecticides was begun in the 1930s with pyrethrins – derived from an East African plant source – and it took on new priority with the discovery of the insecticidal properties of DDT in 1939.

Major anti-malaria campaigns were pursued in the United States and Europe in the period after the Second World War, and beginning in the late 1940s and 1950s, numerous eradication efforts were begun in developing countries. These typically involved widespread spraying of houses with DDT and other pesticides, and they led to marked reductions in disease incidence in some countries and many locales.

The most systematic effort was initiated by the World Health Organization in 1957, which set as a goal the eradication of malaria at the national level. Over the succeeding twelve years, the WHO led intensive programs in many countries – though notably not the most heavily malarial countries of sub-Saharan Africa – aimed at eliminating malaria or reducing its transmission. The underlying logic of the approach was not to eliminate mosquitoes but to break the cycle of transmission by preventing mosquitoes from surviving long enough to transmit the disease from infected people to uninfected people. Although spraying of insecticides was seen as a crucial component of the strategy, it was clear that this approach alone could not be sustained indefinitely. Because the parasites were believed capable of living in human hosts for as much as three years after the initial infection, the strategy was to spray areas systematically for three years, as a way to stop the transmission (WHO 1998).

At the same time, the program sought to provide people with effective drugs that would treat or prevent malaria, so that the number of infected people could be reduced, further interrupting the cycle of transmission (WHO 1998).

Although some commentators have blamed the nascent environmental movement for halting malaria eradication efforts by banning the use of DDT and other insecticides (for a particularly virulent attack, see the website <http://www.rachelwaswrong.org/>), a more accurate account is that by 1969, it had become apparent that insecticidal spraying had not succeeded – and indeed could not succeed – in the most heavily endemic areas. Although spraying did in fact reduce or eliminate the burden of malaria, even in these countries, the initial plan had never contemplated indefinite programs of spraying or control. Instead, three years had been expected to suffice. But the emergence of drug resistant malaria parasites and DDT-resistant mosquitoes had enormously complicated the efforts to break the cycle of transmission, and *ex post* it became evident that it was impractical to eliminate malaria in simultaneously everywhere – and that without

simultaneity, the disease would be reintroduced into countries where it had been previously eliminated (WHO 1998).

As a result, the WHO decided to review and re-evaluate its control strategies beginning in the early 1970s. Over the succeeding several decades, malaria control and prevention efforts fell drastically. In some countries, the disease had been effectively controlled, and these countries had little reason to make it a priority. In the countries where it remained endemic, malaria control efforts were often seen as costly and ineffective, and they competed with many other priorities for investment dollars.

In recent years, however, malaria has re-emerged as an area of intense interest for the international community. As noted above, dozens of new initiatives have sought to revive research on malaria (and particularly on vaccine development) as well as control efforts that involve bed nets and drug therapies.

For the purposes of this paper, it is useful to think of four possible strategies that are used for malaria prevention, control, and treatment. These are: the use of long-lasting insecticide treated bed nets (LLIN); indoor residual spraying of pesticides (IRS); drug treatments, either prophylactic or curative; and vaccines. At the moment, drug options are somewhat limited; only a few effective drugs are available, and epidemiologists worry about the development of resistance by the parasites. The vaccine option is also purely illusory at this writing; no vaccine has been shown to be effective and safe in humans, although some progress has evidently been made towards this goal.

Our paper will consider the economics of these various policy approaches as well as their impact on the level of malaria. Our approach is to ask how each of the alternative policies would alter the steady-state equilibrium that the model economy would attain over time. In each steady state, we can compute the proportion of the population that is sick and the malaria death rate, as well as the level of income per capita and the overall utility or welfare of individuals in the economy. We note that a recurring feature of the model is the existence of multiple steady states; the same policy regime may lead to different outcomes in two otherwise identical economies with different initial conditions.

3. Model

The model used in this paper is based heavily on a companion paper, Gollin and Zimmermann (2007). In the model economy of that paper, individuals are born into a world in which malaria is present. Depending on the initial conditions of the model and on the available prevention and control methods, malaria may be eradicated or may persist and spread. In our model, output per person will depend

in part on the health status of the population. If many people are sick, incomes will be negatively affected. Equally, however, economies that are initially rich are likely to be healthy: high incomes lead to good health, and vice versa.

In our model economy, we track the dynamics of the disease, along with production and prices. Malaria affects individual labor productivity levels and increases the mortality rates of the sick. People in the model economy face persistent idiosyncratic shocks to labor income, in addition to the risks from disease. They lack any formal insurance markets, but households can accumulate assets as a form of precautionary savings. In a sense, our model is thus a Huggett (1996) economy with 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.” Specifically, individuals choose whether or not to seek protection from the disease. We assume that there exists some bundle of goods that offer protection from malaria. This deserves some comment, because obviously there is no vaccine available, and measures such as drugs, insecticides, and bed nets all offer limited protection at best. We will consider explicitly the role played by less-than-perfect efficacy of the preventive goods.

In general, the demand for the preventive good will depend on the cost, the efficacy, and the prevalence of the disease (i.e., the proportion of people sick). But note that there is an important externality in this world. In deciding whether or not to purchase the preventive good, individuals do not take into account the effects that their decisions will have on others. This implies that the decentralized equilibrium will in general have an inefficiently low level of the preventive good. We consider a variety of government policies that would seek to remedy the underprovision of preventive goods.

3.1 *Model environment*

Like any macroeconomic model, ours is abstract to the point of caricature. Our model world is populated with a large number of individuals, who are all born identical. They have different life experiences, however, and there is a large amount of *ex post* heterogeneity.

In the model, individuals are born healthy and without any assets. They live an indeterminate number of periods; *i.e.*, in each period, they face a positive probability of dying, but this probability is independent of age. In each period of their lives, they work, save and consume, as is typical of most macroeconomic models.

While they are alive, people face a probability of getting sick with malaria. We model the disease in a fashion that greatly simplifies the disease's etiology and epidemiology. Specifically, we treat infection as a one-time event. In the model economy, individuals who have been infected remain sick and infectious for the remainder of their lives. They suffer from a permanent reduction in labor productivity, and they face an increased probability of dying, relative to those who have never been infected. We recognize that these assumptions are incompatible with the actual epidemiology of malaria, but they greatly simplify the computation of the model, and we will argue that the qualitative results of the model would be robust to a more realistic treatment of the disease, in which people could recover from the disease and could then experience multiple episodes – perhaps of decreasing severity – over their lifetimes.

In the model economy, the disease has a relatively simple epidemiology. An individual faces higher probabilities of infection if the ecology is particularly conducive to the disease and also if there are more people currently infected. This means that infection rates are endogenous to the model.

Individuals hold no assets at birth, but subsequently they accumulate or spend down assets through their lives. These assets can be used in production; the market for asset services is perfectly competitive. However, there is neither borrowing nor lending in the model; nor is there any insurance market. Therefore, individuals in the economy hold precautionary savings to protect themselves from idiosyncratic productivity shocks, to which they are subject. Assets vanish when people die.⁵

Our treatment of disease prevention and control is critical to this paper. We want to consider a number of different prevention and control options. For simplicity, we need to be able to nest these within a single model, even though they may have somewhat different characteristics. As noted above, our approach is to model different treatments as pairings of cost and efficacy. We will then consider public provision and private provision of different treatments.

In terms of the model environment, we adopt the simple formulation of Philipson (2000), in which a lumpy preventive good may be purchased at any point in an individual's life. Just as we assume that sickness is a permanent condition, so also we view "prevention" as a permanent state; *i.e.*, once an individual has purchased the preventive good, he or she is permanently protected from infection. The degree of efficacy, however, may vary. Some treatments (such as bednets) may cost little but have relatively low efficacy in preventing

⁵ 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.

infection. Other treatments (such as vaccination) may cost more but may have higher long-term efficacy.

More traditional epidemiological models include more possible states of health: individuals can be infected and sick or infected but asymptomatic. They may be contagious or not contagious, symptomatic or not symptomatic. Prevention may last a short time or a long time. To include these details in our model would add significantly to the computational intensity, but we do not believe that qualitative results would change. Clearly, however, our quantitative results would differ to some degree.

3.2 *Preferences and endowments*

In our model economy, a representative individual i gains utility from consumption and health status according to the current- period utility function:

$$u(c_{it}) = \ln(\gamma c_{it})$$

with lifetime utility given by: $\sum_{t=0}^{\infty} \beta^t u(c_{it})$. In this formulation, the parameter s_{it} reflects the potential utility cost of malaria, independent of its impact on consumption allocations. In other words, we allow for the possibility that ill health alone affects utility. Thus, we allow the parameter to take on two possible values, 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 allow us to calculate a steady-state level of mere subsistence; this in turn will allow us to pin down the “value of life” for people in the model economy.

Conditional up on being sick or healthy, individuals choose consumption and savings in each period, along with any purchases of preventive goods, subject to a budget constraint. Their income comes from labor plus any dividends they receive from renting out their assets for productive uses.

Individuals earn labor income from supplying one unit of time in each period; since there is no preference for leisure in our model, they supply time inelastically to the labor market. An individual’s actual supply of effective labor units will depend on health status, s_{it} , and also on an idiosyncratic shock, π_{it} , which is realized in each period. This shock 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:

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

As noted above, individuals may choose to protect themselves from malaria through a one-time purchase of prophylaxis. We model this purchase as a basket of consumption goods, with q_j defining the quantity of consumption that must be foregone for a particular method of prophylaxis, j . For a given method, the individual's binary choice to purchase or not purchase protection at period t is denoted by p_t , such that $p_{it} \in \{0,1\}$.

Given this setup, the individual's period budget constraint under prevention method j is given by:

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

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

3.3 Technology

We model the technology side of the economy with an aggregate production function that displays constant returns to scale. Individual effective labor units aggregate to $L_t = \sum_i h_{it} \pi_{it}$, and individual asset holdings aggregate to $K_t = \sum_i k_{it}$.

These are used to produce output Y_t according to the Cobb-Douglas production function:

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

We assume that factor markets are perfectly competitive, so that factor prices are given by marginal products.

3.4 Model dynamics

In the model economy, we will track births and deaths, although we will impose the restriction that population levels will remain constant over time. To achieve this, we will set fertility rates at replacement levels.

As noted above, individuals face a probability of dying in each period. This probability depends only on an individual's health status; no other variables influence it. Thus, let d_h and d_s be the death rates of healthy and sick people, respectively. Let their fertility rate be f .

As in Gollin and Zimmermann (2007), we define N to be the total population, with S and H respectively denoting the fractions sick and healthy. The fraction V are protected with efficacy e_j through prophylaxis, and the fraction actually purchasing protection in each period is denoted by P .

Armed with this notation, we can write the laws of motion that determine the size of each group within the economy.

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:

$$I = ZS^\mu$$

where Z is an index of malaria ecology, and μ is a parameter. 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. 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 fraction sick and the ecology are at 1, this is another steady state.

The dynamics of the model are driven by the law of motion for the aggregate capital stock: $K' = K + Y - C - PqN - d_s K_s - d_h K_h$, where C is the aggregate consumption, and 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; there is a high level of *ex post* heterogeneity in the model based on the idiosyncratic shocks and health shocks.

3.5 *Equilibrium*

As in Gollin and Zimmermann (2007), we define an equilibrium recursively as 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.

We solve the model computationally. For most specifications and parameterizations, there are multiple steady states. There are in particular for two steady states: one that is attained if the economy has high initial health levels and aggregate assets, and another that is attained if the economy has low initial health levels and asset holdings. In the first steady state, no one is infected and no one will ever be infected. In the second, there is the potential for a fraction of the population to be sick. For our purposes in this paper, this second steady state is the more interesting one, as we try to assess the impact of disease prevention policies.

Finally, we note that the model economy – as is typical of models with infectious disease – is characterized by an important externality related to the transmission of disease. An individual contemplating the decision of whether or not to purchase prevention does not take into account the potential impact of her decision on the infection rates faced by others. As a result, private actors are likely to purchase inefficiently low levels of the preventive good. Thus, there may be a role for the government to subsidize the bundle of preventive goods.

4. Calibration

The calibration of this model is discussed in detail in Gollin and Zimmermann (2007). We will only summarize this process here. A number of the parameter

values for the model are taken as standard. The discount factor β is 0.95, corresponding to a one-year time period. The risk aversion parameter is taken as $\rho = 1$. We compute a value for the parameter $\gamma = 11.3$, to reflect a reasonable value for life. We assume that a person infected with malaria experiences a permanent subsequent loss of 10 percent in his or her efficiency units of labor. The idiosyncratic shocks to productivity are set at a magnitude of 0.224 (following Domeij and Heathcote 2004), with a transition matrix of:

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

We use a capital share on the aggregate production technology of 0.36, in keeping with standard practice in the literature.

Death rates are taken to be 0.075 for sick people (i.e., those infected with malaria) and 0.015 for healthy people. The fertility rate is chosen to ensure a stable population in equilibrium.

Costs of various prevention and treatment methods, and estimates of their efficacy, are taken from Johnson (2007), who calculates lifetime protection costs for insecticide-treated bednets, long-lasting insecticidal bednets, indoor residual spraying, and vaccines. Johnson does not report a full range of efficacy figures, but she cites numbers of theoretical efficacy levels of different treatments.

Finally, we have the parameters Z and μ for the infection rate process. We set $Z = 0.7$ and $\mu = 0.122$ for this exercise. See our other paper for details.

5. Experiments and Results

Using the model economy, we conduct a series of experiments to explore the potential impact of a set of different malaria control and prevention efforts. Specifically, we explore the use of insecticide-treated bednets, long-lasting insecticidal bednets, indoor residual spraying, and vaccines. Each of these interventions is modeled as having a cost and a level of efficacy. We also explore the impact of making these treatments available free, which can be seen as an evaluation of a program to subsidize these treatments.

Our estimates of cost and efficacy are based loosely on Johnson (2007) and are presented in Table 1. Johnson's cost estimates draw on a set of assumptions about the fixed and variable costs of reaching a large population, and they include full costs of delivery. They also incorporate assumptions about the duration of a

single treatment. For example, ITNs need to be re-treated with insecticide every two to three years, and the nets themselves need to be replaced periodically. Similarly, a vaccine may need to be renewed every five to ten years. The marginal cost of vaccination rises steeply as people in more remote areas are included in the coverage; by contrast, the marginal cost of a bednet does not increase very much.

In the data, the most expensive control method appears to be a vaccine, primarily due to the rising marginal cost of delivery. Regular bednets are also relatively costly, due to the need to re-treat with insecticide at fairly regular intervals. By contrast, IRS and LLIN are relatively low-cost methods of control.

Levels of efficacy also vary across methods of control. Vaccines turn out to have relatively low efficacy, which is surprising at first glance. However, the current generation of vaccines is not expected to provide anything close to full protection from all strains of malaria. The parasite is genetically more complex than any organism for which vaccines have yet been developed, and its complexity also implies that it will be difficult to find vaccines that are widely effective. The parasite appears to display a high degree of polymorphism in its surface proteins, so that a vaccine which is effective against one strain of malaria may well fail to protect against other strains. To date, the most effective vaccine trial has achieved only about 30 percent efficacy, and there are no immediate prospects of vaccines that will achieve significantly higher levels of efficacy (Sutherland 2007).

Efficacy of other methods is thought to be somewhat higher. Bed nets are widely believed to achieve protection of 70-80 percent by reducing nighttime exposure to mosquitoes. Since most malarial mosquitoes are nocturnal, bed nets provide a relatively high degree of protection when used conscientiously. However, many observers have suggested that the nets are not often used as intended. Moreover, some have suggested that widespread use of bed nets can lead to a reduction in the mosquitoes' nocturnal habits; it is not yet clear whether the reported efficacy rates are sustainable.

The efficacy of pesticide spraying, by contrast, is not much in question. Spraying works by reducing the density of the mosquitoes that are vectors of transmission for the disease. Spraying is typically targeted to indoor sleeping areas in affected areas, and the sprays have a relatively long-lasting insecticidal effect. Although mosquitoes tend to develop resistance to specific insecticides, there are a number of effective sprays available, and there is some possibility that they can be used in rotation.

For the purposes of this paper, we take the cost and efficacy estimates as accurate. The goal of our paper is not, however, to choose among the four

alternative control methods. If it were, we could quickly eliminate the dominated methods – such as vaccines – that have low efficacy and high costs. But instead, our goal in this paper is to answer the following questions:

- To what extent do people trade off costs of protection against efficacy? What should be the priorities for research: reductions in the cost of protection, or improvements in the efficacy of available methods?
- To what extent is private cost a deterrent to protection? What would be the economic impact of government subsidies that effectively bring down the costs of protection to zero?
- How large are the disease’s likely impacts on economic outcomes (e.g., national income), given the availability of current control methods?
- Do different control methods rank differently in different disease ecologies? Are different methods preferable in the malaria periphery, relative to the centers of disease impact?

To address these questions, we run our calibrated model once for each combination of malaria ecology and treatment type. As a benchmark, we consider an economy in which no effective protection is available. In other words, we model this economy as one in which there is no bundle of protective goods available – perhaps a “natural state.”

As shown in Table 2, this benchmark leads to high rates of infection, with the actual level depending on the ecology. With a moderate malaria ecology ($Z = 0.5$), approximately 85 percent of the population is infected, while with a relatively high malaria ecology ($Z = 0.7$), the infection rate increases to 90 percent. At $Z = 0.9$, 92 percent of the population is infected; essentially, it is only newborns who are healthy, because they have not yet been exposed to the disease.

Output in the benchmark economy is dependent on the malaria ecology. The economy with relatively moderate malaria ecology has an output per person approximately 6 percent higher than the economy with the worst malaria ecology. Consumption is about 10 percent higher in the less malarial economy, reflecting the fact that people live longer and accumulate higher levels of asset holdings, which they are then able to consume. (People try to eat all their assets before they die.) Unsurprisingly, utility is correspondingly highest in the economy with the most favorable malaria ecology.

With the available treatment alternatives, our model finds limited impacts. The costs of these prevention and control methods is low enough – even though the costs represent a significant fraction of income, at roughly one dollar per day

– that essentially everyone in the model is able to accumulate enough to buy protection early in life. In effect, they are able to buy the preventive good in their first year of life, when they are healthy. (In equilibrium, they have no reason to buy the good if they are already sick.) This result is consistent with our earlier findings (Gollin and Zimmermann 2007) that people will spend up to two years' average income on preventive goods, if they are perceived as effective.

Nevertheless, in our model, large fractions of people end up sick, because the treatments are not fully efficacious. Thus, the fraction sick equates with the fraction “sick but protected” in each ecology.

The second panel of Table 2 reports the results of our analysis of an economy in which ITN bed nets are available but costly. The first three columns of the table show the outcomes for an economy in which people are responsible for buying the nets themselves, whereas the second three columns assume that the nets are fully subsidized; i.e., that the price is zero, with the full cost borne by an international actor outside the model.

This fraction who are sick is directly related to the efficacy of the different methods of prevention and control. When vaccines are the available method of prevention, over 80 percent of the people end up sick. The most successful form of prevention, spraying, nevertheless leaves 68 percent of the population sick, even in the mildest disease ecology. All the rest become sick eventually, and since we have modeled the disease as having a lifelong impact, this means that the steady-state has a higher sickness rate than the efficacy of the treatment might suggest.

Why do people continue to buy prevention, even when it inevitably fails? The answer is that it prolongs an individual's period of healthy life, on average. By postponing the day at which the individual becomes sick, it allows an increase in expected duration of life, and thus in lifetime income. In addition, people derive utility from living, so they are willing to spend on goods that help them achieve longer life.

Note that this may hold even when average annual consumption is lower. For example, consider the first panel of Table 2. In this panel, we find that annual consumption is lower with all four treatments than with no treatment at all, in part because almost all individuals are bearing the cost of the preventive good. Indeed, annual output per person is lower in three of the four treatments – with IRS the lone exception – because individuals generally accumulate lower assets as a result of their purchases of the preventive goods. But because people have protection, they live longer. The longer average life expectancy implies that individuals prefer to live in this world, so steady-state utility is always higher in the world where the preventive goods are available.

Of the four possible preventive goods, IRS always delivers the highest utility, which is unsurprising given the assumed advantage in efficacy and cost. What is perhaps surprising is that the increases in output and consumption per person are relatively small, even with this method of prevention available. Relative to the steady-state equilibrium with no available prevention, there is actually a loss in consumption per worker when relatively ineffective control methods are made available in areas with moderate malaria ecology. There are utility gains, but these would not necessarily be measured by examinations of economic indicators. This is an important point to make: some of the benefit of malaria control consists of improvements in health conditions and life expectancy, which will not necessarily be included in national income and product accounts.

The production gains of malaria prevention and control show up more clearly in areas with worse malaria ecology. With a malaria ecology index of 0.7, utility is higher under all treatments than without treatment, and average annual consumption is higher for all the treatments except vaccines. With a malaria ecology of 0.9, which would correspond to extremely severe malarial conditions, all four treatments produce an increase in income per capita, by a minimum of 1.5 percent (with vaccines) to about 7.5 percent (with IRS). These are not huge increases, but they are not inconsistent with the kinds of figures estimated by Weil and Lester (2007) and some other recent analyses.

Next, consider the following experiment. Suppose that instead of having to pay for each of the four methods of prevention and control, individuals had access to these treatments at zero cost. This would correspond, for example, to a situation in which an external donor (such as the World Health Organization or the Gates Foundation) supplied the preventive goods free to anyone who seeks them. For simplicity, we assume that the distribution and management of these goods is also costless. We ignore the cost to these outside actors, and we also ignore any efficiency losses that might be expected to arise (for example, we do not allow bed nets to be misused as fishing nets or wedding veils, as some critics have described).

The question we ask is: how much could this free distribution be expected to improve welfare in the model economy? For simplicity, we focus on the single case of intermediate malaria ecology; *i.e.*, $z = 0.7$. We find that in this case, the effects on consumption are modest. They are also, not surprisingly, inversely related to the cost of the prevention methods. Thus, the biggest gains would come for vaccines: since these are expensive, it would have a relatively large impact to provide these for free, rather than to make people pay for them. Specifically, free provision of the vaccine would increase average consumption by over 2 percent, and average output by just under 1 percent. There would be little effect on the proportions of people sick: in our model, essentially everyone

in the population would already be taking advantage of the vaccination opportunity if it were available, so the increased well-being comes entirely from the financial effects of providing individuals with the subsidy. Essentially the same result holds for the other prevention methods that we consider: because the entire population already makes use of them, subsidized provision would have a financial effect rather than a health effect on the economy. These effects are likely to be small, at least relative to some of the claimed impacts of malaria on income levels (e.g., Gallup and Sachs 2001, Sachs and Malaney 2002).

6. Conclusions

Our model offers a stylized view of malaria's impact and of potential control methods. We welcome further studies on the cost and efficacy of different treatment methods, and we also recognize the desirability of a model that offers a more nuanced depiction of the disease itself; e.g., one in which people can recover from the disease, so that infection is not a life-long condition. However, we do not feel that adding such detail would change the central conclusion of the model, which is that individuals are willing to pay significant amounts – large fractions of one year's income, for example – to protect themselves, even imperfectly, from malaria.

We believe that the current prevalence of malaria in developing countries does not reflect the complete lack of prevention methods or their high costs, so much as it reflects the imperfect effectiveness of those methods that are available. While reductions in cost will be welcome, as will foreign assistance in the fight against malaria, the larger problem is the lack of strongly effective methods of prevention and control.

Our model suggests that, if malaria imposes large costs on individuals, their willingness to pay for prevention methods will be high – implying that uptake rates for effective measures will also be high. If our estimates are wrong, and if malaria does not impose large costs on individuals, then willingness to pay may not be great – but neither will be the impact of the disease.

In the search for new methods of prevention, control, and cure, our model also suggests that low efficacy can have the undesirable effect of making some individuals worse off. Introducing a vaccine, for example, that is costly and has low efficacy will create the potential for reducing the welfare of those individuals who pay for the vaccine and nevertheless contract the disease. In effect, these are individuals who assess *ex ante* that the expected utility of a vaccine is positive. They are making a gamble of a kind that is common in economics; but some will lose by spending money on ineffective prevention methods. The real-world corollary to this point is that governments and donors need to be aware that

individuals can actually be made worse off from well-intentioned efforts to offer new methods of prophylaxis and treatment.

Finally, we note that our model, like most disease models, has difficulty in evaluating the life-saving impact of health care measures. Although we can compare incomes and consumption levels across steady states, these are poor measures of welfare, since there are large differences in the rates of people who are sick and who die. We do compute the utility that individuals gain from reduced probabilities of death, but in the real world, reductions in sickness and death from improved health care measures have an intrinsic moral value that is not easily captured in the model.

7. References

- Carter, Richard and Kamini N. Mendis. 2002. Evolutionary and historical aspects of the burden of malaria. *Clinical Microbiology Reviews* 15(4): 564-594.
- Chakraborty, Shankha, Chris Papageorgiou, and Fidel Perez-Sebastian. 2007. Diseases and Development. Manuscript, Department of Economics, University of Oregon.
- Gallup, John Luke and Jeffrey D. Sachs. 2001. The economic burden of malaria. *American Journal of Tropical Medical Hygiene* 64(1,2)S: 1-11.
- 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.
- Gollin, Douglas and Christian Zimmermann. 2007. Malaria: Disease impacts and long-run income differences. Manuscript: Department of Economics, University of Connecticut.
- Hammer, Jeffrey S. 1993. "The economics of malaria control." *World Bank Research Observer* 8(1): 1-22.
- 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.
- Narasimhan, Vasant and Amir Attaran. 2003. Roll back malaria? The scarcity of aid for international malaria control. *Malaria Journal* 2(8), April 15.
- Papageorgiou, Chris, Shankha Chakraborty, and Fidel Perez-Sebastian. 2005. Diseases and Development. Departmental Working Papers 2005-12, Department of Economics, Louisiana State University.
- Sachs, Jeffrey and Pia Malaney. 2002. The economic and social burden of malaria. *Nature* (415: 7 February): 680-85.
- Sutherland, Colin. 2007. A challenge to the development of malaria vaccines: polymorphic target antigens. *PLoS Med* 4(3): e116 doi: 10.1371/journal.Pmed.0040116.
- World Health Organization. 1998. Fifty Years of the World Health Organization in the Western Pacific Region: Report of the Regional Director to the Regional Committee for the Western Pacific. Online as of August 1, 2007 at: http://www.wpro.who.int/NR/rdonlyres/04A4F011-833E-4603-9442-D04C837DCF1A/0/RC49_03.pdf.

Tables and Figures

Table 1: Cost and efficacy estimates for different approaches to malaria prevention and control.

Method	Lifetime Cost per Person Protected	Efficacy (percent reduction in infection probability)
Insecticide-treated bednets (ITN)	\$45	0.70
Long-lasting insecticide nets (LLIN)	\$20-30	0.70
Residual spraying (IRS)	\$16-32	0.80
Vaccine	\$50-75	0.50

Source: Johnson (2007)

Table 2: Results and experiments*Malaria ecology index 0.5*

Treatment	Sick	Sick/Prot.	K	Y	C
None	0.85	0.00	3.40	1.47	1.27
ITN	0.75	0.75	3.12	1.43	1.18
LLIN	0.75	0.75	3.17	1.44	1.19
IRS	0.68	0.68	3.43	1.49	1.23
Vaccine	0.81	0.81	2.99	1.40	1.14

Malaria ecology index 0.7

Treatment	Sick	Sick/Prot.	K	Y	C
none	0.90	0.00	2.96	1.39	1.16
ITN	0.75	0.75	3.17	1.44	1.18
LLIN	0.75	0.75	3.17	1.44	1.19
IRS	0.68	0.68	3.43	1.49	1.23
Vaccine	0.81	0.81	2.97	1.40	1.14

Malaria ecology index 0.9

Treatment	Sick	Sick/Prot.	K	Y	C
none	0.92	0.00	2.93	1.38	1.15
ITN	0.75	0.75	3.17	1.44	1.18
LLIN	0.75	0.75	3.17	1.44	1.19
IRS	0.68	0.68	3.43	1.49	1.23
Vaccine	0.81	0.81	2.99	1.41	1.15

Table 3: Results and experiments*Malaria ecology index 0.7, free protection*

Treatment	Sick	Sick/Prot.	K	Y	C
None	0.90	0.00	2.96	1.39	1.16
ITN	0.75	0.75	3.17	1.44	1.19
LLIN	0.75	0.75	3.17	1.44	1.19
IRS	0.68	0.68	3.46	1.49	1.24
Vaccine	0.81	0.81	3.03	1.41	1.17