Journal of Health Economics 31 (2012) 658–675

Contents lists available at SciVerse ScienceDirect

Journal of Health Economics

journal homepage: www.elsevier.com/locate/econbase

Diversify or focus? Spending to combat infectious diseases when budgets are tight☆

Soren T. Anderson[a,b,1], Ramanan Laxminarayan[c,d,2], Stephen W. Salant[e,f,*]

[a] Department of Economics, Michigan State University, East Lansing, MI 48824, USA
[b] NBER, USA
[c] Center for Disease Dynamics, Economics & Policy, 1616 P Street NW, Washington, DC 20036, USA
[d] Princeton University, USA
[e] Department of Economics, University of Michigan, 611 Tappan Street, Ann Arbor, MI 48109, USA
[f] Resources for the Future, USA

ARTICLE INFO

Article history:
Received 10 January 2011
Received in revised form 13 March 2012
Accepted 4 May 2012
Available online 18 May 2012

JEL classification:
I10
I18
C61

Keywords:
Dynamic programming
Dynamic optimization
Infectious disease
SIS dynamics

ABSTRACT

We consider a health authority seeking to allocate annual budgets optimally over time to minimize the discounted social cost of infection(s) evolving in a finite set of groups. This optimization problem is challenging since the standard SIS epidemiological model describing the spread of the disease contains a nonconvexity. Neither optimal control nor standard discrete-time dynamic programming can be used to identify the optimal policy. We modify the standard dynamic programming algorithm and show how familiar, elementary arguments can be used to reach conclusions about the optimal policy. We show that under certain conditions it is optimal to focus the entire annual budget on one group at a time rather than divide it among several groups, as is often done in practice. We also show that under certain conditions it remains optimal to focus on one group when faced with a wealth constraint instead of an annual budget.

1. Introduction

Infectious disease remains an important cause of poor health in less developed countries, despite improvements in hygiene, sanitation, vaccination, and access to treatment (Lopez et al., 2006). Even in highly developed countries, diseases such as influenza and HIV/AIDS remain public health challenges. Although vaccines are available for some diseases, treating individuals who are already sick and trying to cure them is the only available intervention for many important diseases such as malaria, cholera, gonorrhea, and tuberculosis. Treatment, though not as effective as vaccination, is therefore an important tool in preventing the spread of infectious disease.

Unfortunately, limited budgets often constrain cost-effective treatment efforts. Unprecedented resources have been devoted to combating HIV, for example, yet the four million people in treatment represent less than 40% of those living with the disease. In Zambia, a country with one of the best-funded malaria control programs in sub-Saharan Africa, only 13% of children with malaria receive effective treatment. Such problems are pervasive. The World Health Organization (WHO), the Global Fund to Fight AIDS, TB and Malaria (GFATM), and individual ministries of health all operate under limited budgets. As a result, different countries or regions are often competing for the same funds. In the case of GFATM, for example, individual countries apply for money to fund disease control projects; some countries receive donated funds, while others do not.

When faced with such constraints and multiple infected populations, these agencies typically allocate funding in proportion to the number of people infected. The GFATM explicitly gives priority to low-income countries with high disease burden. This strategy

☆ For helpful comments and suggestions, we thank two anonymous referees, Arjada Bardhi, Mark Gersovitz, Natalia Lazzati, Michael Murray, and seminar participants at Bates College, CIREQ, and the World Congress of Environmental and Resource Economists in Montreal. We also thank Itamar Megiddo and Zachary Stangebye for valuable research assistance. Any remaining errors are our own.
* Corresponding author at: Department of Economics, University of Michigan, 611 Tappan Street, Ann Arbor, MI 48109, USA. Tel.: +1 734 764 2370; fax: +1 734 764 2769.
E-mail addresses: sta@msu.edu (S.T. Anderson), ramanan@cddep.org (R. Laxminarayan), ssalant@umich.edu (S.W. Salant).
[1] Tel.: +1 517 355 0286; fax: +1 517 432 1068.
[2] Tel.: +1 202 328 5085; fax: +1 202 939 3460.

0167-6296/$ – see front matter
http://dx.doi.org/10.1016/j.jhealeco.2012.05.001

seems equitable, but does it minimize the overall burden of disease? That is, to make the most of their limited budgets, should health authorities devote most treatment to groups with many infected people? Or should authorities focus on groups with many susceptible people? Or, as standard economic intuition might suggest, should they divide their budgets by equating the marginal impact of the last dollar of treatment spent on each group? Since the stated objective of these agencies is to reduce the burden of disease, the question of whether to diversify or focus is central to their missions.

In this paper we consider a health authority allocating treatment between two or more distinct groups to minimize the discounted social cost of infections over a finite time horizon (or, in the limit, an infinite one). The groups could represent populations in schools, hospitals, towns, cities, or countries, or could represent genders or ethnic groups, provided they are distinguishable. The infection in each group is assumed to spread according to the conventional susceptible-infected-susceptible (SIS) epidemiological model. People suffering from the types of infections described by this model do not become immune once they recover but instead become susceptible to re-infection. This type of model is appropriate for certain bacterial agent diseases, such as gonorrhea, meningitis, plague, and streptococcal sore throat, and protozoan agent diseases, such as malaria and sleeping sickness (Hethcote, 2008). The health authority can treat infected people at a constant marginal cost, but treatment each year is limited by a fixed annual budget.

We focus initially on the case where the infection cannot spread between groups. We show under tight budgets and other plausible assumptions that it is never optimal to divide the annual budget between the groups. Instead, the health authority should devote its entire budget in every period to just one group. Further, when there are two identical groups whose only difference is their starting level of infection, it is optimal to treat the sick in the group with the larger number of uninfected people. Since this group starts out healthier and gets all the treatment in the first period, it remains healthier in the subsequent period. Thus, as long as the budget remains insufficient to treat *every* infected individual, it is optimal to focus on the infecteds in the healthier group, period after period, to the complete neglect of sick people in the other group. We later extend our analysis to the case where the disease can spread *between* groups. Although focusing on one group at a time remains optimal in most circumstances, we do identify situations where it is optimal to divide the annual budget between groups.

Our findings that it is often optimal to focus run counter to both conventional practice and standard economic intuition. These results derive from the way an infection spreads, as described in the SIS model of disease. New infections arise from healthy people interacting with the sick. Thus, treating one sick person not only cures that single individual some percentage of the time but then also prevents healthy people from becoming infected at a later date. If there are many such healthy people, then spending the money required to treat one sick person prevents much disease. If many people are already sick, however, then treating one sick person prevents disease in fewer healthy people, and the treated individual is more likely to become sick again. The health authority in effect faces dynamic increasing returns to treatment in each group: the greater the number of healthy people, the more effective treating sick people in the group becomes. Thus, when presented with multiple infected groups and a limited budget, the health authority should take advantage of increasing returns by devoting its entire budget to a single group. Put differently, given the SIS dynamics, the health authority's cost-minimization problem is concave, leading to a corner solution in every period.

Determining how best to minimize the burden of infectious disease calls for a combination of epidemiological and economic insights, an approach taken in both the economics and the epidemiology literatures. Based on the pioneering work of Revelle (1967), Sanders (1971), and Sethi (1974), a more recent literature has emerged to clarify a number of important issues associated with this dynamic optimization problem (Goldman and Lightwood, 2002; Rowthorn and Brown, 2003; Gersovitz and Hammer, 2004, 2005; Smith et al., 2005a; Herrmann and Gaudet, 2009). None of these articles, however, describes the optimal treatment of multiple populations when the health authority has a limited budget.

Most of the literature minimizes the discounted sum of treatment costs plus the social costs of the infection. As always, the solution to such a "planning problem" is a valuable benchmark, since it identifies what is socially best. Often, however, health authorities in the real world are unable to achieve this first-best outcome. An authority may be charged, for example, with minimizing forgone production or school attendance due to illness but may lack the authority to tax or borrow. It then has no choice but to live within its annual budget. Indeed, governmental ministries of health may be prohibited by law from borrowing, as are entities such as the GFATM. In our base case, we assume that no one will lend to this health authority despite its promise to repay the loan out of its future annual budgets – perhaps because the health authority cannot precommit to repaying the loan in the future. To show that our results do not depend on this assumption, however, we also examine the less plausible case where the health authority can borrow against future budgets.[3]

The dynamic increasing returns to treatment inherent in the standard epidemiological model (SIS) make deriving the optimal treatment policy difficult, even for a single population. This nonconvexity in the planner's cost-minimization problem has haunted the literature from the outset. In an early paper that uses dynamic programming, Sanders (1971) concludes that treatment should always be set to zero or the maximum possible level, but Sethi (1974), analyzing the same problem using optimal control methods, concludes to the contrary that optimal treatment is always interior except in transitional phases at the beginning and end of the program. More recently, Gersovitz and Hammer (2004) recognize that they cannot prove analytically that the solutions to their necessary conditions are optimal, since, as they note, the standard sufficiency conditions fail in the presence of the nonconvexity (pp. 10 and 26). They instead rely on numerical simulations to argue that their solutions are likely optimal.

Goldman and Lightwood (2002) show how the nonconvexity inherent in the SIS dynamics can sometimes be handled. In the absence of diminishing returns to treatment, the Hamiltonian in their optimal control problem is linear in the control and, as they note, "comparisons must be made along all paths satisfying the necessary (or first order) conditions." While they are unable to pin down the single optimal path, they are able to reduce the number of possible paths to just two candidates, and these two candidate paths are very easy to characterize. They skillfully demonstrate this strategy in solving a planning problem involving a single population.

Rowthorn (2006) analyzes both the case of a budget-constrained and the case of a wealth-constrained health authority treat-

[3] This latter case involves the minimization of forgone production due to illness subject to any wealth constraint of the health authority. It therefore includes the special case in which the constraint itself has been set optimally, so that, at the constrained optimum, spending a dollar more in wealth would reduce the social cost of infection by a dollar or less; it also includes the case in which the constraint is so tight that spending a dollar more in wealth would reduce the social cost of infection by more than a dollar.