TUTORIALS IN
OPERATIONS RESEARCH
INFORMS 2009

 | ISBN 978-1-877640-24-7
DOI 10.1287/educ.1090.0063

OR's Next Top Model: Decision Models for Infectious Disease Control

Elisa F. Long
School of Management, Yale University, New Haven, Connecticut 06520,
elisa.long@yale.edu

Margaret L. Brandeau
Department of Management Science and Engineering, Stanford University,
Stanford, California 94305, brandeau@stanford.edu

Abstract Infectious diseases continue to threaten many populations worldwide, and timely control of these epidemics is a global health priority. Mathematical models can be used to estimate the expected costs and benefits of alternative disease prevention, diagnostic, and treatment interventions and, ultimately, to aid policy makers with allocating limited disease-control resources. In this tutorial, we describe a variety of operations research-based models that can be used to analyze such problems. We describe models ranging from simple linear models and decision trees to complex network-based simulations. We also describe Markov models with single or multiple decision epochs, individual microsimulation models, population-level dynamic compartmental models (both deterministic and stochastic), and linear programming models. We illustrate each class of model with a published example, focusing on models developed to assess HIV/AIDS prevention and treatment policies. We also discuss decision and model scope, potentially relevant model outcomes, and possible decision criteria for implementing particular health interventions.

Keywords health policy modeling; mathematical model; cost-effectiveness analysis; HIV/AIDS

1. Introduction

Infectious diseases pose a significant burden to public health. Each year more than 13 million people die from infectious diseases, representing approximately 25% of deaths and 50% of premature deaths worldwide (World Health Organization [45]). Infectious diseases disproportionately afflict developing countries (Table 1) and children (Mathers et al. [30]). Among adults aged 15 to 59 in low- and middle-income countries, human immunodeficiency virus (HIV), tuberculosis (TB), and lower respiratory infections (e.g., pneumonia, bronchitis, and influenza) kill four million people annually, devastating economies and straining fragile healthcare systems (Mathers et al. [30], World Health Organization [46]). Children in developing countries are particularly vulnerable to contracting infectious diseases: seven of the top 10 causes of death among children under age 14 are infectious diseases (Mathers et al. [30]). More than six million children die each year from infectious diseases, with malaria and diarrheal diseases accounting for three million childhood deaths alone (Mathers et al. [30]). Other common infectious diseases, including hepatitis, influenza, cholera, dengue fever, typhoid, and yellow fever, continue to plague many countries, potentially leading to surges in new cases, such as the recent cholera epidemic in Zimbabwe that infected 80,000 people and claimed 4,000 lives as of February 2009 (World Health Organization [48]).

Table 1. Ten leading causes of death worldwide in 2004 (World Health Organization [46]).

Low-income countries	Deaths (millions)	High-income countries	Deaths (millions)
Lower respiratory infections	**2.94**	Coronary heart disease	1.33
Coronary heart disease	2.47	Stroke and cerebrovascular disease	0.76
Diarrheal diseases	**1.81**	Trachea, bronchus, lung cancers	0.48
HIV/AIDS	**1.51**	**Lower respiratory infections**	**0.31**
Stroke and cerebrovascular disease	1.48	Chronic obstructive pulmonary disease	0.29
Chronic obstructive pulmonary disease	0.94	Alzheimer's and dementia	0.28
Tuberculosis	**0.91**	Colon and rectum cancers	0.27
Neonatal infections	0.90	Diabetes mellitus	0.22
Malaria	**0.86**	Breast cancer	0.16
Low birth weight	0.84	Stomach cancer	0.14

Note. Infectious diseases are highlighted in bold.

Because of their enormous health impact, control of infectious diseases is a key global public health priority. A variety of interventions, including vaccination, prevention programs, and treatment regimens, can reduce new infections and disease-related mortality. For example, for HIV, a range of prevention programs—including voluntary counseling and testing, condom promotion, prevention of mother-to-child transmission, education campaigns, needle-exchange and harm-reduction programs for injection drug users, and more recently, male circumcision (Kahn et al. [23])—have been shown to be effective and cost-effective methods of preventing new infections. Antiretroviral therapy for HIV-infected individuals has generated millions of incremental life-years since its inception in 1996 and, because it reduces infectivity in treated patients (Cohen et al. [10], Erb et al. [16]), has likely also prevented many new HIV infections. As another example, control of vector-borne malaria with mosquito bed nets, indoor residual spraying, and prophylactic drugs has prevented millions of malaria infections (World Health Organization [47]). Although a malaria vaccine is not yet available, results of a recent clinical trial suggest that such a vaccine may become available in the not-so-distant future (Bejon et al. [4]). Despite such progress, millions of people continue to become infected and die from HIV and malaria each year. Similar progress and challenges exist for numerous other infectious diseases.

Although most infectious disease deaths are preventable, resources for disease control are limited. Thus, when determining the appropriate response to each epidemic, policy makers at the local, national, and international levels must determine the appropriate allocation of scarce public health resources. To make informed decisions, policy makers need to know the likely costs and health impacts of competing disease-control programs. However, estimation of the costs and health impacts of programs that aim to control an infectious disease is complicated by the fact that infectious diseases are *dynamic* (epidemics evolve over time), *nonlinear* (the rate of new infection is approximately a function of the number of people who are infective multiplied by the number of people who are susceptible), and *stochastic* (many factors, such as behavioral and biological factors that influence the transmission and progression of a disease vary across individuals and over time).

This is where operations research (OR) can play an important role: OR-based models can quantify the likely economic costs and health benefits of disease-control programs, providing important information about the potential benefits of implementing alternative interventions, as well as the consequences of failing to do so. OR-based models can also determine—explicitly or implicitly—the most efficient allocation of limited disease-control resources.

This tutorial describes a range of models that can be used to quantify the potential effects of infectious disease-control programs and thus support decision making about which

programs to implement. We first discuss decision scope and framing, and the consequent implications for modeling (§2). We next describe relevant health and economic outcomes (§3). We then describe a variety of types of models used to support decision making for infectious disease control, ranging from simple linear models to complex stochastic dynamic models (§4). We illustrate each class of model with published examples that were developed to assess HIV prevention and treatment policies. We conclude in §5.

2. Decision Scope and Framing

OR-based models can be used to inform a variety of decisions about infectious disease control. The scope of the decision often dictates the appropriate type of model to use. Before committing to a particular mathematical framework, it is useful to consider a number of factors. Who are the decision makers? What are their objectives? What are the goals of the analysis? What information is available about the disease? Answering these and other questions helps frame the decision problem to determine the type(s) of models that may be most suitable for addressing the problem.

2.1. Patient-Level Decisions

Some decisions about infectious disease control, such as decisions about treatment strategies, may be specific to individual patients or specific subsets of patients. In such a case, it is important to know the likely effect of the intervention on individual patients. Because disease diagnosis, progression rates, and treatment response may differ for each patient, a model that captures individual variability may be useful when analyzing such a problem. For example, among HIV-infected individuals, the increase in viral load in the bloodstream and the rate of decline in immune system function depends on a number of patient-specific factors (e.g., patient age, comorbidities, and time since infection) as well as chance. To capture the variable effect of HIV treatment on patient outcomes, constructing a patient-level model is essential. Models of individualized HIV disease management often include a complex array of health states and transitions between health states, but typically ignore or approximate the effects of disease transmission on other individuals in the population (Freedberg et al. [17], Paltiel et al. [37], Sanders et al. [41]). Such a framework is most appropriate for informing patient-level decisions (e.g., when should an individual begin a treatment regimen, how often should a patient receive viral load monitoring, etc.).

2.2. Population-Level Decisions

A higher-level decision relates to disease control in an entire population, such as decisions about broad-based prevention and vaccination programs. For these decisions, it is important to know the likely populationwide consequences of alternative programs. This societal perspective aggregates across the entire population the net health benefits and economic costs associated with each disease-control measure. For example, when considering rapid national scale-up of HIV treatment or the introduction of universal HIV screening, a national health ministry would want to estimate the future course of the HIV epidemic under different program scenarios, and the associated populationwide program costs and healthcare costs.

Instead of considering the best course of action for a single individual, the decision maker who is considering broad-based interventions must balance the interests of individuals with those of the entire population. This is especially critical for eradicating or diminishing infectious disease epidemics. For example, many parents choose not to vaccinate their children against measles because they fear potential adverse side effects, but mathematical models suggest that a minimum measles vaccination coverage of 90%–95% is needed to achieve herd immunity in a population (Wallinga et al. [44]). Although the optimal individual decision may be to decline vaccination when most others in the population are vaccinated, the optimal decision from a societal perspective is to promote or require mass childhood vaccination.

Population-level decisions are inherently more complex than patient-level decisions because each individual may be acting independently, and the decision maker needs to balance these competing interests. The advantages of using a societal perspective instead of a patient-level perspective are amplified when modeling an infectious disease because of the potential externalities. An individual with heart disease who fails take medication impacts no one else, whereas a patient infected with a communicable disease such as TB who poorly adheres to treatment can infect other people with a more severe drug-resistant strain. The primary disadvantage of societal-based models is their simplification of individual variability and differences in the natural course of the disease and treatment effectiveness.

2.3. Resource Allocation

The goal of modeling infectious disease epidemics is not only to estimate future economic and health outcomes under different control scenarios but also to determine the optimal allocation of limited resources between competing interventions. In its most general form, one can think of this as a problem of optimizing health benefits subject to a constraint on available resources and perhaps other constraints. The appropriate resource allocation criterion will depend on a number of considerations, including budget constraints, political will, reimbursement precedents, and demands from patient advocates. Many tools from operations research, such as decision analysis, optimization, and control theory, have been applied to resource allocation decisions in health care to quantify the trade-offs of implementing alternative strategies.

Cost-effectiveness analysis and cost-benefit analysis are standard tools for comparing the costs and benefits of two or more medical interventions (Gold et al. [20]). The incremental cost-effectiveness ratio (ICER) is calculated as the marginal cost of an intervention divided by the marginal benefit. It measures how much additional "bang for the buck" could be achieved by switching from one intervention to another. This can be written as

$$\text{ICER} = \frac{Costs_{\text{With Intervention}} - Costs_{\text{Without Intervention}}}{Benefits_{\text{With Intervention}} - Benefits_{\text{Without Intervention}}}.$$

Before determining the optimal disease-control policy, decision makers must choose the appropriate decision criteria (that is, they must choose how to select programs based on the cost-effectiveness information) and appropriate constraints (these may include budget constraints, limitations on allowable allocations of resources, etc.). We now discuss three possible ways of framing the resource allocation problem.

2.3.1. Cost-Effectiveness Threshold. One possible decision criterion is to define a cost-effectiveness threshold and then implement any health intervention below this threshold; such an allocation makes the implicit assumption that funds not spent on a particular program (one that is above the cost-effectiveness threshold) could be spent elsewhere on a cost-effective program (below the threshold) that targets another population or another disease. In the United States and Western Europe, a threshold of $50,000 or $100,000 per quality-adjusted life year (QALY) gained is often cited, based on the historical cost effectiveness of kidney dialysis, a commonly implemented intervention (Ubel et al. [43]). However, many critics argue that this threshold was arbitrarily chosen and has not altered over the past 20 years (Ubel et al. [43], Braithwaite et al. [8]). Indeed, many commonly accepted health interventions cost more than $100,000 per QALY gained; for example, many women in the United States continue to receive annual Pap smears for detecting cervical cancer at an estimated cost of $675,000 per QALY gained (Eddy [14]). The World Health Organization has set guidelines for cost effectiveness in developing countries: interventions are deemed cost effective if they cost less than three times the annual gross domestic product (GDP) per capita and highly cost effective if they cost less than the GDP per capita (Murray et al.

[35]). This has prompted a call by researchers to reassess the definition of "cost effective" in many settings (Ubel et al. [43], Braithwaite et al. [8]).

2.3.2. Budget Limitation. An alternative decision rule is to rank a list of interventions by their incremental cost-effectiveness ratio, from most favorable to least favorable. The decision maker can go down the list and select interventions until a predefined budget is exhausted. Although in principle this avoids the undesirable task of determining the monetary value of a year of life, in practice it is often difficult to determine the appropriate budget. With multiple payers, including patients, insurance companies, public sector payers (e.g., Medicare and Medicaid), it is unclear if such a budget threshold even exists, let alone what value the threshold should take. If medical interventions offer substantial gains in health benefits, perhaps it is wise to increase the funds available for such interventions? These considerations make this decision rule difficult to implement in practice.

2.3.3. Equity Constraints. Policy makers may want to consider objectives other than economic efficiency, such as equity among population groups or programs. For example, state health departments may wish to allocate disease prevention resources among counties in proportion to their population, without regard to disease prevalence in each county. A proportional allocation of resources among population subgroups (or among disease-control programs) may achieve equity, but often this is at the expense of program efficiency (Kaplan and Merson [24]). Policy makers could consider these competing objectives and allocate resources in a way that balances equity and efficiency.

3. Relevant Outcomes

Good decision making about infectious disease control requires careful estimates of relevant health and economic outcomes (Drummond et al. [12]). In this section we describe the different types of outcomes that may be appropriate to include in an OR-based model of infectious disease control.

3.1. Economic Outcomes

The economic cost of a medical intervention includes the direct cost of the intervention itself (e.g., prescription drug cost, vaccine cost, prevention program cost, etc.) as well as all indirect costs of related health care (e.g., hospital visits, ancillary services, etc.). If the medical intervention lasts for a finite amount of time, the model should also capture future healthcare costs occurring after the completion of the program. One caveat to this approach, however, stems from the ongoing debate among economists regarding whether future costs should include medical costs unrelated to the intervention under consideration (Garber and Phelps [18], Meltzer [32], Lee [28], Garber and Phelps [19], Meltzer [33]). The debate centers on whether the cost of an intervention leading to reduced mortality should include future healthcare consumption or if this is implicitly accounted for with increased survival.

For example, suppose an individual infected with latent TB receives isoniazid prophylaxis for six months to prevent active TB disease. Assuming future costs are included, a model of this intervention should include the cost of the prophylaxis, all costs of related clinical visits, adherence services, etc., and healthcare costs after six months, to capture the future costs associated with preventing active TB disease. Last, all costs should be discounted to the present, typically at an annual rate of 3% to 8% (Gold et al. [20]).

3.2. Health Outcomes

The counterpart to an intervention's cost is the expected health benefit it generates in the population. Outcomes specific to a particular disease or intervention—such as number of infections averted or number of children who die—are often relevant when comparing interventions for controlling a particular disease, but such outcomes fail to provide a universal standard when comparing multiple diseases.

To compare across diseases and interventions, a more standard metric such as life years or quality-adjusted life years (QALYs) is needed. Life years account for variations in mortality; QALYs additionally adjust each time period based on the patient's quality of life. In a mathematical model, QALYs are computed by integrating the amount of time spent in different health states, multiplied by the quality-of-life (QOL) factor for each health state. For example, an individual who lives for three months in perfect health (QOL 1.0), then spends six months in a moderately sick health state with QOL 0.8, and then another three months in a state with QOL 0.4 before dying would experience 0.75 total QALYs ($0.25 \times 1.0 + 0.5 \times 0.8 + 0.25 \times 0.4 = 0.75$). As with costs, health benefits should also be discounted to the present at the same rate. If this important step is missed, then the optimal strategy in the first year would be to defer implementing the program until the second year because it costs less but attains the same health benefit, and this would continue ad infinitum.

3.3. Epidemic Outcomes

Epidemic outcomes may also be relevant for models of infectious disease control. Disease prevalence refers to the fraction of the population (or number of individuals) infected with the disease at a particular point in time. Incidence measures the number of new cases occurring over a period of time (and the incidence rate is the number of new cases per person per unit time). For example, 1.1 million Americans are estimated to be infected with HIV, resulting in an HIV prevalence of 0.3%, and HIV incidence is approximately 56,300 new cases per year (an incidence rate of 0.0002 cases per person-year) (Centers for Disease Control and Prevention [9]).

Another epidemic measure that is sometimes relevant is the basic reproduction number, (R_0), which is defined as the average number of secondary infections caused by a typical infected individual during his entire period of infectiousness, in a completely susceptible population, absent any interventions. R_0 represents the strength of an infectious disease at sustaining itself in the population: when $R_0 < 1$ the disease will die out in the long run; when $R_0 > 1$ the disease will remain endemic in the population. In general, higher values of R_0 correspond to diseases that are more difficult to eliminate. R_0 depends on the average duration of infectiousness, the probability of transmitting the infection per contact, and the rate of new contacts between infected and uninfected groups. Some commonly referenced values of R_0 are as follows: influenza (2.0–3.0), HIV (2.0–5.0), severe acute respiratory syndrome (SARS) (2.0–5.0), smallpox (5.0–7.0), and measles (12.0–18.0). R_0 should be interpreted with caution, however, because it is a theoretical outcome derived from mathematical analyses. By the time a new epidemic raises public health concerns, the population is no longer entirely susceptible (some individuals are either infected or immune), and many formal interventions (e.g., school closures, quarantines, prophylactic antibiotics) and informal interventions (e.g., voluntary social distancing) may be in place. However, R_0 can often support other epidemiologic evidence and can provide a basis for vaccination policies and other infection control measures.

4. Classes of Models

We now describe the types of models that are frequently used to analyze infectious disease control decisions. For ease of comparison between model types, we focus on one disease, HIV/AIDS, and describe the types of decisions considered and models used. This section is roughly organized according to model complexity, with the simplest types of models discussed first and the most complex models discussed last. Figure 1 provides a schematic diagram of each type of model.

FIGURE 1. Schematic diagram of models.

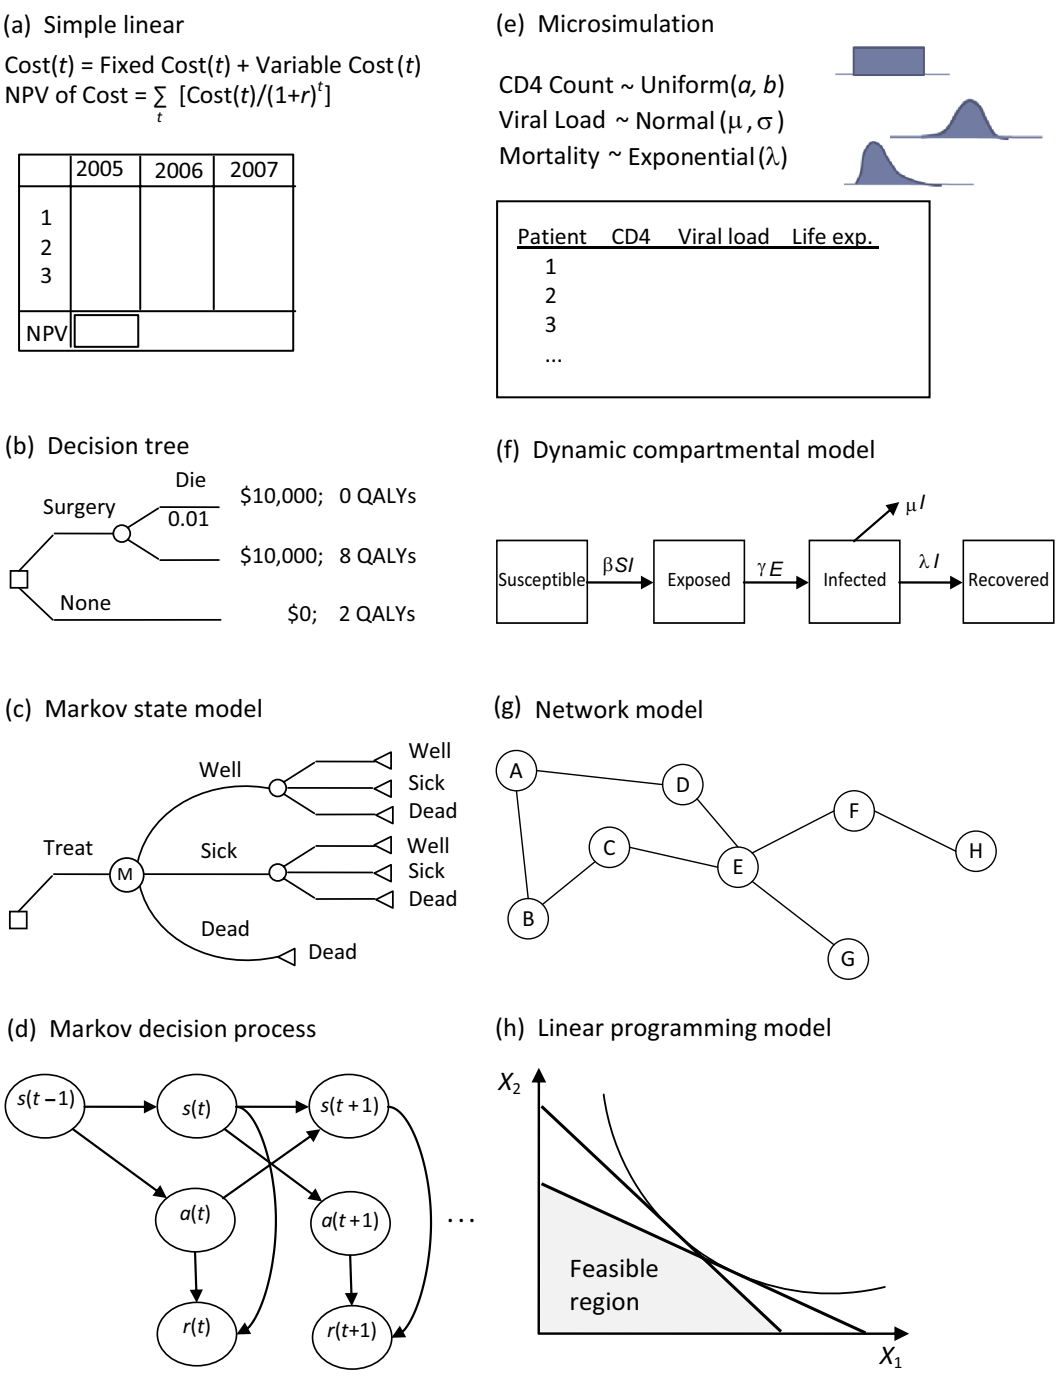

4.1. Simple Linear Models

Simple linear models (Figure 1(a)) can provide transparent analysis of the estimated costs and benefits of disease interventions. For example, Kahn et al. [23] used a deterministic linear model to estimate the cost effectiveness of adult male circumcision (which reduces the probability of HIV infection in heterosexual men) in sub-Saharan Africa. The authors calculated the cost of the intervention as the sum of the direct cost of the intervention and

the indirect cost of potential adverse events, such as post-surgery infection:

$$Costs = (number\ circumcised \times unit\ cost) + (number\ of\ adverse\ events \times adverse\ event\ cost).$$

The effectiveness of the intervention is measured in number of HIV infections averted:

$$\begin{aligned} Effectiveness = {} & number\ circumcised \times (1 - HIV\ prevalence) \\ & \times incidence\ rate \times net\ protective\ effect \times 20\ years \times epidemic\ multiplier, \end{aligned}$$

where

$$net\ protective\ effect = 1 - [(1 - protective\ effect) \times (1 + risk\ compensation)],$$

which accounts for the benefits of male circumcision in reducing the probability of HIV transmission in men (*protective effect*), and the potential increase in risky behavior because of a false sense of protection (*risk compensation*). The *epidemic multiplier* parameter approximates the number of secondary infections in women that are prevented by circumcising their male partners.

The analysis showed that under most reasonable sets of assumptions, adult male circumcision will not only reduce the spread of HIV but will also reduce costs (Kahn et al. [23]). Thus, it is a cost-saving HIV prevention program. Although this modeling framework does not fully capture the changed HIV epidemic dynamics that might accrue from such a program, nor other potentially relevant effects (e.g., reductions in other sexually transmitted diseases because of circumcision), the analysis provides a useful "order of magnitude" calculation and shows that the proposed prevention program is almost certain to be highly cost effective. Simple linear models of this type can provide enormous value by identifying key insights and encouraging the development of more detailed models in the future, if necessary and appropriate.

4.2. Decision Trees

Both simple linear models and decision trees typically provide deterministic outcomes (e.g., expected cost and health benefits of a decision). However, linear models ignore chance events, whereas decision trees explicitly specify a probability distribution over a set of possible outcomes. Decision trees (or probability trees) can be used to model a sequence of uncertain events and the resulting outcomes (Figure 1(b)). The cost effectiveness of alternative intervention strategies can be calculated by comparing the expected costs and benefits for each strategy (i.e., decision). Decision trees have been used to model both patient-level and policy-level decisions.

Bertolli et al. [6] developed a patient-level decision model for estimating the benefits of alternative strategies for preventing mother-to-child HIV transmission and the resulting mortality among children under age five. HIV transmission can occur prenatally or during childbirth. Such transmission can be reduced by administering maternal antiretroviral therapy. An infant who is born uninfected to an HIV-infected mother can still contract the virus later if the mother breastfeeds. However, if the mother abstains from breastfeeding, the child has a higher chance of dying before age five because of the lack of critical antibodies acquired from breast milk and the risk of substituting unclean drinking water. Figure 2 presents an illustrative decision tree for this situation. The decision tree captures the uncertainties associated with these possible events, and their effects on HIV infection, other infections, and mortality.

Decision trees can also be used to compare HIV prevention policies at the population level. Nakchbandi et al. [36] compared mandatory versus voluntary HIV screening for pregnant women in the United States. The model includes a cohort of pregnant women who may elect to receive or refuse prenatal care in the presence of a mandatory HIV screening program. An

FIGURE 2. Example of a decision tree for estimating the effectiveness of preventing mother-to-child HIV transmission.

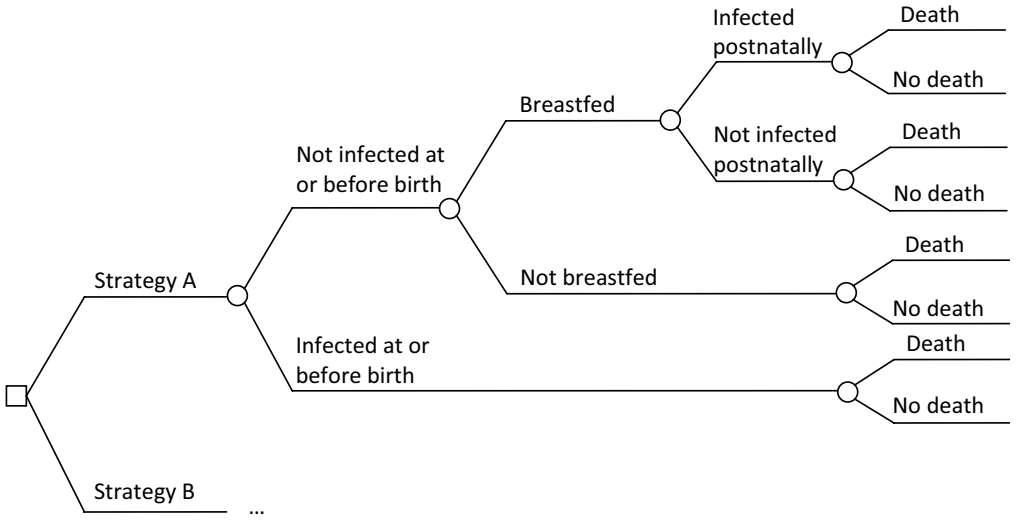

alternative to mandatory screening is voluntary screening, which may reach fewer women but could reduce the number deterred from receiving care. Once again, the decision analysis framework allows the authors to include variations in the conditional probability of each event, given a prior sequence of events. The authors measured the net population benefits in terms of total HIV-infected and dead infants under different screening policies.

4.3. Markov Models

Decision trees are useful for modeling instantaneous or one-time probabilistic events, but they fail to capture the effect of spending time in a particular state. Markov models (Figure 1(c)) achieve this by incorporating a finite-state Markov chain and the corresponding transition probabilities. Health states can correspond to a particular disease stage, or the presence of treatment or some other disease-related status (e.g., circumcision status, screening status, vaccination status, etc.). Patients can transition between states according to the natural progression of the disease, underlying biological or behavioral factors, or due to the presence of some control measure. Markov models integrate well with cost-effectiveness analysis: it is straightforward to sum up the total time spent in each health state and multiply this by the appropriate quality-of-life factor; costs for each health state are calculated similarly, and the direct intervention cost is determined with a reward (or toll) for transitioning between states.

Sanders et al. [41] developed and applied a Markov model to evaluate the cost effectiveness of HIV screening in the United States. The model follows a cohort of patients over their lifetime and simulates HIV infection status, viral load, CD4 count (a measure of the immune system's strength), treatment status, and deaths from HIV and other causes. The model incorporates behavioral, clinical, and epidemiologic data, and has been used to estimate the health benefits and costs of expanded HIV screening in the United States. The findings from this study and from a microsimulation-based study by Paltiel et al. [37] (described in §4.5) prompted the U.S. Centers for Disease Control and Prevention to recommend universal HIV screening of adults and adolescents in all healthcare settings (Branson et al. [7]).

4.4. Markov Decision Processes

Markov models of the type described in §4.3 typically assume that only one decision is made, at the start of the model. From that point forward, a sequence of events unfolds according

to state-dependent transition probabilities. A Markov decision process (MDP) extends a simple Markov model by allowing for multiple decision points over time (Figure 1(d)). Each decision epoch is defined by a set of possible health states, S, a set of possible actions, A, a probability matrix, P_a, and a reward matrix, R_a.

$$P_a(s, s') = \Pr(s_{t+1} = s' \mid s_t = s, a_t = a).$$

$P_a(s, s')$ is the probability of transitioning from state s to s' if action a is taken at time t, and $R_a(s, s')$ is the immediate reward associated with transitioning from state s to s' if action a is taken at time t.

The solution to an MDP is calculated recursively and is defined by an optimal policy $\boldsymbol{\pi}$ and a value function $\mathbf{V}$. These are written as follows:

$$\pi(s) = \arg\max_a \sum_{s'} P_a(s, s')V(s'),$$
$$V(s) = R(s) + \beta \sum_{s'} P_{\pi(s)}(s, s')V(s'),$$

where β is the one-period discount factor. One disadvantage of using MDPs in cost-effectiveness analysis is that they typically optimize in one dimension (minimizing costs or maximizing health benefits). Once the optimal policy is determined, the other metric is also calculated; then the costs and benefits of optimal policies with different objective functions can be compared using cost-effectiveness analysis.

Shechter et al. [42] developed a Markov decision process to determine the optimal time to initiate antiretroviral therapy in HIV-infected patients. Initiating treatment too early can lead to adverse side effects, drug toxicities, and exhaustion of available treatment options, whereas initiating treatment too late can lead to the development of AIDS-related symptoms and possibly death. For this problem, such a sequential decision model is useful because patients typically see their clinician at regular time intervals, and physicians must decide when to initiate treatment based on established guidelines and individual patient characteristics.

4.5. Microsimulation

Most decision tree and Markov models provide deterministic estimates of the expected costs and benefits of disease interventions. These models can be extended to include a microsimulation (Figure 1(e)), where a hypothetical cohort of patients can be sent through a model, with each patient generating one sample path. The mean and variance can then be calculated for all patients, in a manner similar to calculating the average effects in a real clinical trial. A microsimulation ("first-order simulation") represents individual patient *variability*, whereas a traditional Monte Carlo simulation ("second-order simulation") accounts for parameter *uncertainty*. As an example, consider rolling a "fair" die and receiving \$1 × the number rolled. A deterministic model would calculate the expected payoff (\$3.50), a microsimulation model would repeatedly roll the die and calculate the mean and variance across all rolls, and a Monte Carlo simulation would replace the payoffs ($\$1, \$2, \ldots, \$6$) with probability distributions and repeatedly sample each distribution simultaneously.

Bendavid et al. [5] simulated a cohort of 100,000 hypothetical HIV-infected patients to estimate the benefits of alternative disease monitoring strategies in patients in sub-Saharan Africa. For each patient, the model simulates, on a month-by-month basis, age, HIV treatment status, CD4 cell count, HIV viral load, the development of opportunistic infections, and possible medication toxicity. As another example, the CEPAC (Cost Effectiveness of Preventing AIDS Complications) model (Freedberg et al. [17], Paltiel et al. [37]), which is designed to evaluate the effects of screening, treatment, and other programs for HIV-infected and uninfected patients, simulates HIV disease progression in infected individuals

over time, and the corresponding rate of immune system (CD4 T-cell count) decline, change in HIV blood plasma viral load, rate of antiretroviral therapy failure, and opportunistic infection development. Microsimulation models are useful for characterizing the potentially vast differences in disease progression, treatment response, mortality, and quality of life among individuals. However, as with Markov models, many microsimulation models fail to appropriately capture disease transmission in the population. Two other classes of models—dynamic compartmental models (§4.6) and network models (§4.7)—are better suited for projecting the future epidemic.

4.6. Dynamic Compartmental Models

Dynamic compartmental models (Figure 1(f)) are used to project the evolution of an epidemic over time. In such a model, the population is divided into a set of mutually exclusive, collectively exhaustive compartments. Transitions of individuals between compartments over time are modeled according to a system of nonlinear difference or differential equations. Compartmental models vary in size and complexity, as well as the underlying mathematical dynamics.

4.6.1. Simple Deterministic (Analytical) Models. In 1927, Kermack and McKendrick [25] wrote a pioneering article discussing the application of nonlinear dynamic systems to epidemic control. In 1979, Anderson and May (Anderson and May [2], May and Anderson [31]) applied this theory to many modern infectious diseases. The underlying idea behind a deterministic epidemic model is to characterize the epidemic by the number of susceptible and infected individuals over time. This is represented by a system of nonlinear differential equations. Deterministic compartmental models with few disease states can often be solved analytically.

One simple model, known as the SIR model (susceptible, infected, recovered), is as follows:

$$\frac{dS}{dt} = -\beta S(t)I(t),$$
$$\frac{dI}{dt} = \beta S(t)I(t) - \gamma I(t),$$
$$\frac{dR}{dt} = \gamma I(t).$$

$S(t)$, $I(t)$, and $R(t)$ correspond to the number of susceptible, infected, and recovered individuals in the population at time t, respectively. The infection rate is denoted by β, and the recovery rate is denoted by γ. The recovery rate can be interpreted as the inverse of the average duration of infectiousness, $1/\gamma$. The basic reproduction number can be calculated analytically:

$$R_0 = \frac{\beta S(0)}{\gamma}.$$

The term R_0 is proportional to the force of new infections in the population, $\beta S(0)$, and the average duration of infectiousness, $1/\gamma$.

In addition to solving for R_0, one can use Lyapunov's indirect method to determine the stability criteria for the steady-state equilibria of the linearized system. For example, Long et al. [29] determined the disease-free and endemic equilibria of a coepidemic model of two infectious diseases. Such analyses are useful for calculating the long-run effect of different control measures on epidemic outcomes and showing the conditions under which an epidemic could theoretically be eradicated.

4.6.2. Simple Deterministic (Numerical) Models. As the number of disease states increases, a compartmental model can more realistically capture the subtleties of disease transmission and progression. At the same time, however, the model can become analytically intractable. In this case, numerically solving the system of differential equations is the best alternative. After specifying initial conditions for each equation, the system can be estimated by using a numerical solution technique such as a Runge-Kutta method for approximating solutions of ordinary differential equations.

Zaric et al. [49, 50] developed a deterministic, dynamic compartmental model to estimate the cost effectiveness of methadone maintenance treatment for injection drug users. The model captures the benefits of methadone to individuals in the treatment program and, more critically, to their sexual and needle-sharing partners. The ability to capture the effects of an intervention (in particular, health effects related to disease transmission) for all individuals in the population, and not just the effects for those who receive the intervention, is a significant benefit of this class of models. Whereas Markov models often assume a fixed incidence rate after an intervention is implemented, compartmental models capture the dynamic evolution of disease incidence over time as the number of susceptible and infected individuals change. Moreover, compartmental models can capture the different costs and benefits accruing to individuals in each different population compartment.

4.6.3. Stochastic Models. Some compartmental models are based on stochastic differential equations, where the underlying system is based on a diffusion process with the disease states and time both treated as continuous variables. With a stochastic SIR model, the states $S(t)$, $I(t)$, and $R(t)$ are random variables, and a unique sample path can be drawn on each iteration. The rate of change in infected people over time is represented by the following stochastic differential equation:

$$\frac{dI(t)}{dt} = \mu(I(t)) + \sigma(I(t))\frac{dW(t)}{dt},$$

where

$$\mu(I(t)) = \beta I(t)S(t) - \gamma I(t),$$
$$\sigma(I(t)) = \sqrt{\beta I(t)S(t) + \gamma I(t)},$$
$$W(t+\Delta t) - W(t) \sim \text{Normal}(0, \Delta t).$$

$W(t)$ is a Wiener process with stationary, independent time increments. A more detailed discussion of stochastic compartmental models can be found in the article by Allen and Burgin [1]. Stochastic epidemic models are particularly useful when random variations in model parameters (e.g., infectivity rate, number of contacts, mortality rate) are important, or if the population under consideration is small and/or heterogeneous (e.g., an sexually transmitted disease outbreak among adolescents). On the other hand, deterministic compartmental models are useful when considering large populations, where model parameters for an *average* individual are sufficient. Stochastic models can become computationally intensive, because multiple simulation runs are required to estimate average epidemic outcomes.

4.7. Network Models

Perhaps the most computationally intensive technique for modeling the effects of an infectious disease-control program entails modeling each individual in the population and his or her partnerships. This is known as a network model (Figure 1(g)). In the network representation of such a model, nodes represent individuals, and links represent contacts between those individuals (for example, sexual partnerships or needle sharing). As the population size

increases, the number of possible contacts grows exponentially. Many network models are therefore limited to isolated or small populations. Each node can contain information about an individual's characteristics such as age, gender, health state, and risk behavior profile. Sophisticated network models can include geospatial coordinates in addition to temporal information. Such models are especially relevant for diseases with geographic dispersion, such as the SARS outbreak in 2003. A recent article suggests that modeling and understanding the complicated network dynamics underlying sexual partnership formation and dissolution is a key step in controlling the HIV epidemic (Koopman [27]).

To evaluate the effects of programs that aim to reduce the number of concurrent sexual partnerships in sub-Saharan Africa (and thus reduce the spread of HIV), Enns et al. [15] developed a microsimulation model of a population of sexually active males and females. The model incorporates a dynamic network model of sexual partnerships (including partnership formation and dissolution) overlaid with a detailed model of HIV transmission and disease progression in each individual. The model includes both spousal and nonspousal partnerships; spousal partnerships are characterized by longer duration but higher chance of infection transmission than nonspousal partnerships. In addition to partnerships, the model simulates each individual's age, HIV infection status, treatment status, and CD4 cell count. The model assigns a cost and quality of life to each health state, and thus can estimate the cost effectiveness of different levels and types of concurrent partnership reduction that may be achieved by a concurrency reduction program. Network models of this type can realistically simulate important details of disease contact networks, but can be complex to create and validate, and are usually data and computation intensive.

4.8. Linear Programming Models

Linear programming models can be applied to optimize a linear function (e.g., health benefits) by allowing the decision variables (e.g., investment in a particular intervention) to vary, subject to a set of constraints (e.g., budget, healthcare capacity, or equity constraints) (Figure 1(h)). The objective function can depend on baseline disease incidence and the level of investment in each intervention (i.e., the intervention's production function). A linear objective function may be a simple calculation based on the expected gain in health benefits assuming different portfolios of interventions are implemented. More complex nonlinear programming may include an underlying disease model, such as a dynamic compartmental model, to capture the effects of disease transmission on health outcomes.

In 2001, the Institute of Medicine commissioned a report to develop guidelines for allocating HIV prevention resources in the United States (Committee on HIV Prevention Strategies in the United States [11]). The authors presented a linear programming framework for maximizing the number of HIV infections averted by deciding how to allocate resources among risk groups (men who have sex with men, injection drug users, and high-risk heterosexuals) in all 50 states. The model estimated the baseline incidence in each risk group, the maximum number of people reached within each population, the program's effectiveness at preventing new infections, and the per-person cost of each program. Using the existing available budget, the authors estimated the amount by which an efficient allocation could increase averted HIV infections. Other examples of linear programming applications for allocating HIV resources have also been documented (Earnshaw et al. [13], Richter et al. [40]).

5. Discussion

Infectious diseases impose a significant burden on public health. OR models can play an important role in informing—and improving—decisions about how to allocate scarce infection-control resources. A variety of models, ranging from simple linear calculations to complex dynamic models, are useful for modeling the effects of possible infectious

disease control interventions. Development of such a model may include concepts from areas such as health outcomes analysis (Drummond et al. [12]), cost-effectiveness analysis (Gold et al. [20]), decision analysis (Raiffa [39]), Markov models (Howard [22]), Markov decision processes (Howard [22], Puterman [38]), microsimulation (Mitton et al. [34]), infectious disease modeling (Anderson and May [3]), network modeling (Knoke and Yang [26]), and optimization (Hillier and LIeberman [21]). The appropriate model depends on the decision under consideration, as well as available data. OR's "next top model" is one that helps generate a significant improvement in global health.

Acknowledgments

Margaret Brandeau was supported by a grant from the National Institute on Drug Abuse (R01-DA15612).

References

[1] L. J. S. Allen and A. M. Burgin. Comparison of deterministic and stochastic SIS and SIR models in discrete time. *Mathematical Biosciences* 163:1–33, 2000.

[2] R. M. Anderson and R. M. May. Population biology of infectious diseases: Part I. *Nature* 280(5721):361–367, 1979.

[3] R. M. Anderson and R. M. May. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, Oxford, UK, 1991.

[4] P. Bejon, J. Lusingu, A. Olotu, A. Leach, M. Lievens, J. Vekemans, S. Mshamu, T. Lang, et al. Efficacy of RTS,S/AS01E vaccine against malaria in children 5 to 17 months of age. *New England Journal of Medicine* 359(24):2521–2532, 2008.

[5] E. Bendavid, S. D. Young, D. A. Katzenstein, A. M. Bayoumi, G. D. Sanders, and D. K. Owens. Cost-effectiveness of HIV monitoring strategies in resource-limited settings: A southern African analysis. *Archives of Internal Medicine* 168(17):1910–1918, 2008.

[6] J. Bertolli, D. J. Hu, P. Nieburg, A. Macalalad, and R. J. Simonds. Decision analysis to guide choice of interventions to reduce mother-to-child transmission of HIV. *AIDS* 17(14):2089–2098, 2003.

[7] B. M. Branson, H. H. Handsfield, M. A. Lampe, R. S. Janssen, A. W. Taylor, S. B. Lyss, and J. E. Clark. Revised recommendations for HIV testing of adults, adolescents, and pregnant women in health-care settings. *Centers for Disease Control Morbidity and Mortality Weekly Report, Recommendations and Reports* 55(RR-14):1–17, 2006.

[8] R. S. Braithwaite, D. O. Meltzer, J. T. King, Jr., D. Leslie, and M. S. Roberts. What does the value of modern medicine say about the $50,000 per quality-adjusted life-year decision rule? *Medical Care* 46(4):349–356, 2008.

[9] Centers for Disease Control and Prevention. Estimates of new HIV infections in the United States. Accessed December 15, 2008, http://www.cdc.gov/hiv/topics/surveillance/resources/factsheets/incidence.htm, 2008.

[10] M. S. Cohen, C. Gay, A. D. Kashuba, S. Blower, and L. Paxton. Narrative review: Antiretroviral therapy to prevent the sexual transmission of HIV-1. *Annals of Internal Medicine* 146(8): 591–601, 2007.

[11] Committee on HIV Prevention Strategies in the United States. *No Time To Lose: Getting More from HIV Prevention*. M. S. Ruiz, A. R. Gable, E. H. Kaplan, M. A. Stoto, H. V. Fineberg, and J. Trussell, eds. Division of Health Promotion and Disease Prevention, Institute of Medicine, National Academy Press, Washington, DC, 2001.

[12] M. F. Drummond, B. J. O'Brien, G. L. Stoddart, and G. W. Torrance. *Methods for the Evaluation of Health Care Programmes*, 2nd ed. Oxford University Press, Oxford, UK, 1997.

[13] S. R. Earnshaw, K. Hicks, A. Richter, and A. Honeycutt. A linear programming model for allocating HIV prevention funds with state agencies: A pilot study. *Health Care Management Science* 10(3):239–252, 2007.

[14] D. M. Eddy. Screening for cervical cancer. *Annals of Internal Medicine* 113(3):214–226, 1990.

[15] E. A. Enns, T. Igeme, E. Bendavid, and M. L. Brandeau. Setting concurrency reduction targets for HIV prevention. Working paper, Stanford University, Stanford, CA, 2009.

[16] P. Erb, M. Battegay, W. Zimmerli, M. Rickenbach, and M. Egger. Effect of antiretroviral therapy on viral load, CD4 cell count, and progression to acquired immunodeficiency syndrome in a community human immunodeficiency virus-infected cohort. Swiss HIV Cohort Study. *Archives of Internal Medicine* 160(8):1134–1140, 2000.

[17] K. A. Freedberg, E. Losina, M. C. Weinstein, A. D. Paltiel, C. J. Cohen, G. R. Seage, D. E. Craven, et al. The cost effectiveness of combination antiretroviral therapy for HIV disease. *New England Journal of Medicine* 344(11):824–831, 2001.

[18] A. M. Garber and C. E. Phelps. Economic foundations of cost-effectiveness analysis. *Journal of Health Economics* 16(1):1–31, 1997.

[19] A. M. Garber and C. E. Phelps. Future costs and the future of cost-effectiveness analysis. *Journal of Health Economics* 27(4):819–821, 2008.

[20] M. R. Gold, J. E. Siegel, L. B. Russell, and M. C. Weinstein, eds. *Cost-Effectiveness in Health and Medicine*. Oxford University Press, New York, 1996.

[21] F. S. Hillier and G. J. Lieberman. *Introduction to Mathematical Programming*, 2nd ed. McGraw-Hill, New York, 1995.

[22] R. A. Howard. *Dynamic Programming and Markov Processes*. MIT Press, Cambridge, MA, 1960.

[23] J. G. Kahn, E. Marseille, and B. Auvert. Cost-effectiveness of male circumcision for HIV prevention in a South African setting. *PLoS Medicine* 3(12):e517, 2006.

[24] E. H. Kaplan and M. H. Merson. Allocating HIV-prevention resources: Balancing efficiency and equity. *American Journal of Public Health* 92(12):1905–1907, 2002.

[25] W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London Series A* 115(772):700–721, 1927.

[26] D. Knoke and S. Yang. *Social Network Analysis*, 2nd ed. Sage Publications, Thousand Oaks, CA, 2007.

[27] J. Koopman. Modeling infection transmission. *Annual Review of Public Health* 25:303–326, 2004.

[28] R. H. Lee. Future costs in cost effectiveness analysis. *Journal of Health Economics* 27(4): 809–818, 2008.

[29] E. F. Long, N. K. Vaidya, and M. L. Brandeau. Controlling co-epidemics: Analysis of HIV and tuberculosis infection dynamics. *Operations Research* 56(6):1366–1381, 2008.

[30] C. D. Mathers, A. D. Lopez, and C. J. L. Murray. The burden of disease and mortality by condition: Data, methods, and results for 2001. A. D. Lopez, C. D. Mathers, M. Ezzati, D. T. Jamison, and C. J. L. Murray, eds. *Global Burden of Disease and Risk Factors*. Oxford University Press, New York, 45–93, 2006.

[31] R. M. May and R. M. Anderson. Population biology of infectious diseases: Part II. *Nature* 280(5722):455–461, 1979.

[32] D. Meltzer. Accounting for future costs in medical cost-effectiveness analysis. *Journal of Health Economics* 16(1):33–64, 1997.

[33] D. Meltzer. Response to "Future costs and the future of cost-effectiveness analysis." *Journal of Health Economics* 27(4):822–825, 2008.

[34] L. Mitton, H. Sutherland, and M. Weeks, eds. *Microsimulation Modelling for Policy Analysis: Challenges and Innovations*. Cambridge University Press, Cambridge, UK, 2000.

[35] C. J. Murray, D. B. Evans, A. Acharya, and R. M. Baltussen. Development of WHO guidelines on generalized cost-effectiveness analysis. *Health Economics* 9(3):235–251, 2000.

[36] I. A. Nakchbandi, J. C. Longenecker, M. A. Ricksecker, R. A. Latta, C. Healton, and D. G. Smith. A decision analysis of mandatory compared with voluntary HIV testing in pregnant women. *Annals of Internal Medicine* 128(9):760–767, 1998.

[37] A. D. Paltiel, M. C. Weinstein, A. D. Kimmel, G. R. Seage, E. Losina, H. Zhang, K. A. Freedberg, and R. P. Walensky. Expanded screening for HIV in the United States—An analysis of cost-effectiveness. *New England Journal of Medicine* 352(6):586–595, 2005.

[38] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New York, 1994.

[39] H. Raiffa. *Decision Analysis*. McGraw-Hill College Publishing, New York, 1997.

[40] A. Richter, K. A. Hicks, S. R. Earnshaw, and A. A. Honeycutt. Allocating HIV prevention resources: A tool for state and local decision making. *Health Policy* 87(3):342–349, 2008.

[41] G. D. Sanders, A. M. Bayoumi, V. Sundaram, S. P. Bilir, C. P. Neukermans, C. E. Rydzak, L. R. Douglass, et al. Cost-effectiveness of screening for HIV in the era of highly active antiretroviral therapy. *New England Journal of Medicine* 352(6):570–585, 2005.

[42] S. M. Shechter, M. D. Bailey, A. J. Schaefer, and M. S. Roberts. The optimal time to initiate HIV therapy under ordered health states. *Operations Research* 56(1):20–33, 2008.

[43] P. A. Ubel, R. A. Hirth, M. E. Chernew, and A. M. Fendrick. What is the price of life and why doesn't it increase at the rate of inflation? *Archives of Internal Medicine* 163(14):1637–1641, 2003.

[44] J. Wallinga, J. C. Heijne, and M. Kretzschmar. A measles epidemic threshold in a highly vaccinated population. *PLoS Medicine* 2(11):e316, 2005.

[45] World Health Organization. Removing obstacles to healthy development. Accessed October 3, 2008, http://www.who.int/infectious_disease_report. World Health Organization, Geneva, Switzerland, 1999.

[46] World Health Organization. The top 10 causes of death. (Fact Sheet 310). World Health Organization, Geneva, Switzerland, 2008.

[47] World Health Organization. World Malaria Report 2008. World Health Organization, Geneva, Switzerland, 2008.

[48] World Health Organization. Daily cholera update and alerts: February 19, 2009. World Health Organization, Geneva, Switzerland, http://www.who.int/hac/crises/zmb/sitreps/zimbabwe_cholera_19feb2009.pdf, 2009.

[49] G. S. Zaric, P. G. Barnett, and M. L. Brandeau. HIV transmission and the cost-effectiveness of methadone maintenance. *American Journal of Public Health* 90(7):1100–1111, 2000.

[50] G. S. Zaric, P. G. Barnett, and M. L. Brandeau. Methadone maintenance and HIV prevention: A cost-effectiveness analysis. *Management Science* 46(8):1013–1031, 2000.