Contact Tracing to Control Infectious Disease:

When Enough is Enough

Benjamin Armbruster*†   Margaret L. Brandeau*‡

September 25, 2007

Abstract

Contact tracing (also known as partner notification) is a primary means of controlling infectious diseases such as tuberculosis (TB), human immunodeficiency virus (HIV), and sexually transmitted diseases (STDs). However, little work has been done to determine the optimal level of investment in contact tracing. In this paper, we present a methodology for evaluating the appropriate level of investment in contact tracing. We develop and apply a simulation model of contact tracing and the spread of an infectious disease among a network of individuals in order to evaluate the cost and effectiveness of different levels of contact tracing. We show that contact tracing is likely to have diminishing returns to scale in investment: incremental investments in contact tracing yield diminishing reductions in disease prevalence. In conjunction with a cost-effectiveness threshold, we then determine the optimal amount that should be invested in contact tracing. We first assume that the only incremental disease control is contact tracing. We then extend the analysis to consider the optimal allocation of a budget between contact tracing and screening for exogenous infection, and between contact tracing and screening for endogenous infection. We discuss how a simulation model of this type, appropriately tailored, could be used as a policy tool for determining the appropriate level of investment in contact tracing for a specific disease in a specific population. We present an example application to contact tracing for chlamydia control.

Keywords: contact tracing – infectious disease – network – cost-effectiveness

*Department of Management Science and Engineering, Stanford University, Stanford, California 94305-4026, USA

†armbruster@stanford.edu

‡brandeau@stanford.edu

1 Introduction

Studies of infectious disease control efforts — such as screening, vaccination, and contact tracing (also known as partner notification) — often focus on evaluating the effectiveness of such programs. For example, Saretok and Brouwers [1] evaluated the likely epidemic impact of school closures in the event of a potential pandemic influenza outbreak; Bozzette et al. [2] evaluated the effectiveness of different vaccination strategies in mitigating the effects of a smallpox outbreak; and Hethcote [3] evaluated the impact on gonorrhea prevalence of six different control measures. Some analyses of infectious disease control efforts have focused on finding the most effective means of allocating a fixed amount of epidemic control resources across competing interventions and populations, and/or over time. For example, Zaric and Brandeau [4] determined the optimal allocation of a fixed budget among HIV prevention programs at a given point in time and over time [5]; Longini et al. [6] determined the optimal distribution of a limited supply of vaccine in the event of an influenza pandemic; and Halloran et al. [7] performed a similar analysis in the event of a smallpox outbreak. Other analyses have evaluated the minimum cost means of achieving a desired level of epidemic control. For example, Müller [8] determined the cost-minimizing pattern of vaccination among different age groups to achieve a desired epidemic reproduction number, and Revelle et al. [9] determined the cost-minimizing use of treatment, prophylaxis, and BCG (bacille Calmette-Guerin) vaccination to achieve a target prevalence of tuberculosis (TB) in a given time horizon.

None of these analyses considers the optimal level of investment in epidemic control. In some cases, epidemic control programs may have diminishing returns to scale: incremental investments in epidemic control can yield diminishing health benefits. Moreover, funds not spent to control one disease can often be spent to control another disease, or on other public health programs. Thus, investing the maximum available funds to control an infectious disease may not represent a cost-effective use of resources. In this paper we present a methodology for determining the appropriate level of investment in contact tracing. We develop and apply a simulation model of contact tracing and the spread of an infectious disease among a network of individuals. We evaluate the cost and effectiveness of different levels of contact tracing, and show how to determine the appropriate level of investment.

Contact tracing is a primary means of disease control for infectious diseases with low prevalence.

In the US, contact tracing is required for TB [10], recommended for human immunodeficiency virus (HIV) [11], and not uncommon for other sexually transmitted diseases (STDs) [12, 13]. Contact tracing has also been used (and modeled) for severe acute respiratory syndrome (SARS) [14], foot-and-mouth-disease [15], smallpox [16, 17], and avian influenza [18].

Hyman et al. [19] and Armbruster and Brandeau [20] studied contact tracing using differential equation models that assume homogeneous mixing of the population. Kretzschmar [21] reviewed STD models on networks. Müller et al. [22] introduced one of the first models of contact tracing on a network and analyzed a stochastic branching process that approximates it. Subsequent work [15, 23, 24] analyzed similar models using both stochastic simulations and moment closures (also called mean-field approximations). Most of these papers study the effectiveness of contact tracing but do not consider the costs. Armbruster and Brandeau [20] and Wu et al. [18] incorporated costs in their analyses, but considered contact tracing as an all-or-nothing decision, with a fixed level of intensity.

Empirical studies of the cost effectiveness of contact tracing programs have been carried out for diseases such as TB [25, 26], HIV [27, 28], chlamydia [29], syphilis [30], and gonorrhea [31]. These studies all evaluate a single fixed level of contact tracing.

Armbruster and Brandeau [20] presented a theoretical model for determining when (a fixed level of) contact tracing should supplement screening to control an endemic infectious disease. That work was based on a simple compartmental model (an SI model) with homogeneous mixing, and employed a highly stylized representation of the contact tracing process (with contact tracing yielding new identified disease cases at a constant rate, as a function of disease prevalence). A number of studies have shown that analyses based on more realistic models of disease transmission in social networks can yield significantly different projections of disease spread than projections generated by simple compartmental models [21, 32]. In addition, depending on how it is implemented (e.g., how many contacts of any index case are traced; which contacts are traced — and possibly removed — first; and other factors), a fixed level of contact tracing may yield different numbers of new index cases identified, as well as differing impact on the spread of the disease. Therefore, we use simulation to evaluate the effectiveness (and cost) of different levels of contact tracing — and thus to determine the appropriate level of investment in contact tracing.

A number of researchers have analyzed disease transmission in social networks, but often with

little knowledge of the actual network structure. Schneeberger et al. [33] used data on the distribution of the number of sexual partners when constructing networks of sexual contacts among individuals in Britain and Zimbabwe. A few studies have reconstructed various small social networks: for example, Klovdahl et al. [34] mapped the network of sexual contacts of injection drug users (IDUs) and prostitutes in Colorado Springs; Weeks et al. [35] described the social network of IDUs in Hartford, Connecticut; Parker et al. [36] mapped a network of sexual contacts among 154 people in London starting from a 20-year-old HIV-infected man; and Helleringer et al. [37] described the network of sexual contacts in a region of Malawi. Wylie and Jolly [38] were able to construct a larger network of sexual contacts among individuals in Manitoba, Canada by using the reported contacts of chlamydia and gonorrhea cases. For respiratory diseases, the contact network may be easier to find: Saretok and Brouwers [1] used the location of the homes and workplaces of Swedes to model the spread of an avian flu pandemic.

The following section describes our simulation model of contact tracing and the spread of an endemic infectious disease among individuals in a network. In Section 3, we simulate the contact tracing and disease process for different amounts of contact tracing. We show that contact tracing is likely to have diminishing returns to scale: incremental investments in contact tracing yield diminishing reductions in disease prevalence. Using the simulation results in conjunction with a cost-effectiveness threshold, we show how to determine the optimal level of investment in contact tracing. We extend the analysis in Section 4 to consider the optimal allocation of a budget between contact tracing and disease screening. We consider the cases of screening for exogenous infection (infected individuals entering the population) and endogenous infection (infected individuals already in the population). In Section 5 we illustrate the use of the model with an example of contact tracing for chlamydia control. We conclude in Section 6 with discussion of results and directions for future research.

2   Simulation Model

Network Structure   We consider an infectious disease that is endemic in a population of n individuals. We model individuals as nodes on an undirected graph where an edge between nodes i and j indicates that infection transmission can occur between i and j if one these individuals is

infected (we say they are *contacts* of each other). We employ an SIRS epidemic model (described below) with five states: susceptible (S), infected (I), removed (R), susceptible and being traced as a contact (ST), and infected and being traced as a contact (IT).

We used random small-world graphs in our initial simulations. We chose this family of random graphs because it is the only one among the major families of random graphs (the others are Erdos-Renyi, regular, and scale-free) that allows for significant clustering and short paths between pairs of nodes. Watts and Strogatz [32] give examples of these features in many real networks and show that they significantly affect the spread of disease on a network.

To generate the graphs, we started with a cyclic regular graph of n nodes with degree 4 where node i connects to nodes $i \pm 1, \pm 2 \pmod{n}$. For every other pair of nodes (i, j) we created a link independently with probability $1/n$. This process creates a network in which each node has approximately five contacts on average. Figure 1 shows a small example of such a network with its nodes in various states. Our choice of n (500) and average degree (5) is consistent with data from the Colorado Springs study [34], which found a main connected network with 669 individuals and a median of 5.1 risky relationships per person (11.7 relationships per person but of which 29% reported no risky behavior and 27% reported drug sharing without needle use).

Epidemic Model   We employ an SIRS epidemic model in which susceptible individuals become infectious, become removed when they are treated, and finally become susceptible after treatment. We assume that no deaths occur in the population over the simulation time horizon. Figure 2 illustrates the disease states and transitions among them. In our simulation, the sojourn time in each state was exponentially distributed for all states except for states ST and IT, where the sojourn time was a constant.

We assume that the rate of *endogenous* infection (transition from $S \rightarrow I$) of node i (or, more precisely, the individual represented by node i) is proportional to d_i, the number of infected neighbors of node i: specifically, the transition rate is d_i/t_1, where t_1 is a time constant. This stochastic process on a network is called a *contact process*. Individuals can also become infected from *exogenous* sources. This could be through international travel, for example, or by healthy people leaving the system and being replaced by infected immigrants. We assume that the rate at which exogenous infection occurs is given by a constant, η_1 among susceptible individuals and η_2 among

removed individuals. The combined rate at which susceptible individuals become infected is then $(d_i/t_1) + \eta_1$.

To model contact tracing, Eames and Keeling [24] and Kiss et al. [15] extended the contact process so that infected nodes are found and cured at a rate proportional to the number of index case neighbors a node has (in our model, this would be individuals in state R), analogous to the infection process. This model of contact tracing does not allow us to compare different contact tracing budgets. Thus, we use a discrete-event simulation.

When an infected individual seeks treatment for symptoms of the disease (and thus becomes known to the public health system), he or she becomes an *index case*. This corresponds to a transition in disease state from $I \rightarrow R$. We assume that this transition happens at rate $1/t_2$, where t_2 is a time constant. When a new index case occurs, we apply our contact tracing policy to decide (based on only the graph structure and the removed nodes) which nodes to trace. Nodes selected for tracing then transition to state ST or state IT, depending on whether the individual is susceptible or infected, respectively. We assume that contact tracing requires a fixed amount of time, t_4 for state ST and t_5 for state IT.

After tracing is completed, a node in state ST returns to state S, whereas a node in state IT transitions to state R and becomes a new index case. We assume that the contact tracing capacity, K, is expressed in terms of the maximum allowable contact tracing rate: at any point in time, at most K individuals in total can be in states ST or IT. We can think of the capacity K as being proportional to the manpower available for contact tracing.

Contact Tracing Process In contact tracing, every index case is asked to name his or her contacts (graph neighbors who may be infected). Then public health officials seek out these contacts (as time and resources permit) to test whether they are infected and treat them if so. Who to trace is an important tactical decision since the contact tracing capacity limits the number of individuals who can be traced at any point in time.

In our simulation we keep a prioritized list of contacts who have not yet been traced (nodes in state S or I that are neighbors of removed nodes). Every time a new index case is identified, we update this list and decide on additional nodes to trace. For our analyses in Section 3 and Section 4 we assume that each index case names all of his/her contacts; for our example of chlamydia contact

tracing in Section 5, we assume that individuals name only a fraction of their contacts. We let k be the number of contacts we would like to trace each time a new index case arrives. Since the list is prioritized, we trace the k nodes of highest priority, provided we have not exhausted the budget.

Based on a simulation study of the effectiveness of alternative contact tracing strategies [39], we selected the following scheme for prioritizing contacts for tracing. We assign each contact a score intended to reflect the likelihood that the contact is infected (the higher the score, the more likely that a contact is infected). The contact's score is the number of index cases who name that person. We set k, the number of contacts we trace each time a new index case arrives (assuming we still have resources), equal to 5.

Costs and Health Outcomes We compare alternative levels of contact tracing based on the resulting effectiveness and annual costs in steady state. We consider both the cost of the contact tracing and the cost of treating disease cases. We assume a cost of c for treating a case of disease — this cost is incurred each time an individual transitions from disease state I to R — and an annual cost of C for each unit of contact tracing capacity (hence an annual contact tracing cost of KC for K units of capacity).

We consider two measures of contact tracing effectiveness: steady-state disease prevalence, and annual steady-state quality-adjusted life years (QALYs) experienced in the population. The steady-state disease prevalence is a simple measure of the effectiveness of a particular disease control strategy. Following standard cost-effectiveness practice [40], we also measure the total QALYs experienced in the population each year. To do so, we assign quality multipliers (on a scale of 0 to 1, where 0 represents death and 1 represents perfect health) to health states. We assign a quality multiplier of 1 to uninfected individuals (those in states S, ST, and R), and a quality multiplier of $q < 1$ to infected individuals (those in states I and IT). Given a steady-state disease prevalence p, the steady-state number of QALYs experienced in the population during one year is $n(1 - pq)$.

Simulation Runs We used the simulation model to estimate for various cases — e.g., for different levels of the contact tracing budget, for different amounts of contact tracing and screening, etc. — the annual steady-state treatment costs (adding the annual cost of the contact tracing, KC, yields the total cost for a year in steady state) and the steady-state prevalence of the disease

(or, equivalently, the steady-state annual QALYs experienced in the population). To measure the steady state for a particular case, we performed 1600 runs. For each run, we generated a random small-world graph and infected a single random node. Then we simulated the network for five years in one-day time increments (1825 days in total), taking the daily average prevalence (per capita frequency of states I and IT) starting with day 181. (Visual inspection of sample runs indicated that the system is in steady state by this time.) We averaged over all the runs and set our error bars to the 95% confidence intervals.

Table 1 shows the values of all parameters we used in our simulations. These parameters are roughly consistent with the transmission and control of gonorrhea. The speed of disease transmission, t_1, is consistent with unprotected sexual contact every 45 days, and a 50% chance of transmission per unprotected sexual contact [41]. Choosing a 30-day average duration between infection and treatment (t_3) is consistent with all women and 50% of men being asymptomatic, symptoms otherwise appearing after 4 days, and then a 3-day delay in obtaining treatment [41, 42]. Gonorrhea is treated with one dose of antibiotics which costs approximately $50 [43]. We estimated that individuals would refrain from risky behavior for an average of 3 months after treatment. Our contact tracing cost, C, ($120 per case, figuring a week $t_4 = t_5 = 5$ per case and 50 work weeks per year) is similar to that reported by Dasgupta et al. [25]. Quality multipliers for STDs are not well researched [44]. We chose a quality multiplier of .9, which has been used for TB, diabetes, and asymptomatic HIV infection [45].

Table 1: Simulation parameters

Parameter	Value	
n	500	individuals
t_1	90	days
t_2	30	days
t_3	90	days
t_4, t_5	5	days
η_1, η_2	1/9000	new cases/day/person
c	50	USD / case
C	6,000	USD / case / year
q	0.9	

3 How many to trace?

Cost Effectiveness Threshold Choosing the budget for contact tracing is an important strategic decision. Funds not spent to trace a particular disease could be used for tracing other diseases, for other disease control efforts, or even for other public health efforts. Thus, we would like to determine the most "cost-effective" level of investment in contact tracing for a particular disease.

Cost-effectiveness analysis can help policy makers allocate money across different interventions for the same or different diseases [40]. In a typical cost-effectiveness analysis of alternative interventions, the analyst evaluates the total costs and health benefits — usually measured by quality-adjusted life years (QALYs) experienced — for each intervention. The analyst then identifies the non-dominated interventions (a dominated intervention is one that costs more and yields fewer QALYs than another single intervention or than a linear combination of two interventions). The undominated interventions can then be ranked in order of increasing cost to create an efficient frontier of interventions. The incremental cost-effectiveness ratio, defined as the incremental cost per QALY gained, increases as one moves along the efficient frontier. The "best" intervention is the one that has the highest possible incremental cost-effectiveness ratio without exceeding a pre-determined cost-effectiveness threshold α. We will suppose in our analyses that a value for the cost-effectiveness threshold α is known. This value is often determined as an implicit value given by accepted public health/medical practice [46].

In our model, the "alternative interventions" involve alternative levels of investment in contact tracing (or in contact tracing and screening). Because our simulation analyses focus on the steady-state disease prevalence achieved by a given level of contact tracing, we express all costs and health benefits in annual terms. Thus, the cost of each "alternative intervention" equals the annual cost of contact tracing (CK) plus the annual cost to treat cases of the disease. The effectiveness is measured in terms of the resulting steady-state prevalence, or equivalently as the total steady-state number of QALYs experienced in the population per year. Thus, for these analyses, the cost-effectiveness threshold α represents the maximum amount we are willing to pay per year for one more steady-state QALY. We used the value $\alpha = \$50,000$, a value commonly used in cost-effectiveness analyses of health-related interventions [46].

Finding the Optimal Level of Contact Tracing Figure 3(a) shows the steady-state disease prevalence as a function of the contact tracing capacity (K ranging from 0 to 10), using our baseline simulation parameters, and averaged over 1600 runs. The convexity of the curve shows that the effectiveness of contact tracing (in terms of reductions in disease prevalence) decreases with its capacity (i.e., it has diminishing returns to scale): for each incremental increase in the contact tracing capacity, the corresponding reduction in steady-state prevalence diminishes. This makes intuitive sense because as the contact tracing capacity increases and prevalence decreases, we trace more contacts, fewer of whom will be infected. Thus, the probability that the contacts we trace are infected decreases as the contact tracing capacity increases. Our simulation model allows us to quantify this decrease — and thus to evaluate the relative cost effectiveness of different amounts of contact tracing.

Figure 3(b), which is based on the same simulations as Figure 3(a), shows the total annual cost associated with each level of contact tracing (K ranging from 0 to 10), and the total annual (steady-state) QALYs experienced in the population (recall that the population size $n = 500$ is constant). The costs incurred include not only the capacity costs of contact tracing ($C = \$6000$ for each unit of capacity K) but also the costs of treating the disease ($c = \$50$ per case treated). While the contact tracing capacity is the biggest annual expense, disease treatment costs cannot be neglected. Indeed, as contact tracing capacity (K) increases, treatment cost decreases because the prevalence of the disease decreases. Our simulation model allows us to quantify these costs and savings.

The curve connecting the 11 points in Figure 3(b) is concave, reflecting the diminishing effectiveness of contact tracing with incremental investment. The optimal budget is given by the point on the curve where the tangent line has a slope of $1/\alpha$. At this point, we spend \$18,000 per year to maintain a contact tracing capacity of 3 (i.e., the ability to trace 3 people at a time) and we spend approximately \$7,500 per year on treatment of the disease, bringing the total cost to \$25,500. At this point, the incremental cost per reduction in QALYs equals the maximum level we will tolerate, α. Above this point, increases in annual expenditure on contact tracing increase the annual number of QALYs experienced by less than 1 per \$50,000 spent (i.e., less than the ratio $1/\alpha$) and are thus deemed not cost effective.

Sensitivity Analysis Using the simulation model, we can determine how the optimal contact tracing capacity will vary as a function of various disease parameters. We performed one-way sensitivity analysis on the disease transmissibility (i.e., baseline disease incidence, which is proportional to the parameter t_1), the rate at which symptoms develop (and thus the rate at which individuals seek and receive treatment, which is proportional to t_2), the expected length of time a treated individual remains in the Removed state (proportional to the parameter t_3), and the expected length of time between exogenous infections ($1/\eta := 1/\eta_1 = 1/\eta_2$). We varied each parameter to 10% above and below the base case.

For each case, Figure 4 shows the steady-state disease prevalence as a function of the contact tracing capacity. Table 2 shows the effect of changes in these parameters on the optimal capacity (the capacity at which the marginal total cost per QALY gained equals the cost-effectiveness threshold).

Table 2: Sensitivity analysis: Approximate change in contact tracing capacity

	-10% reduction	10% increase
Transmissibility (t_1)	Increases by 1	Decreases by 2
Symptom development (t_2)	Decreases by 2	Increases by 2
Removed time (t_3)	Increases by 1	Decreases by < 1
Exogenous infection ($1/\eta_1$, $1/\eta_2$)	Increases by < 1	Decreases by < 1

For any given contact tracing capacity, the effectiveness of contact tracing (measured by the steady-state prevalence) is affected the most by the rate at which symptoms develop. If that rate is slower than in the base case (and thus there are more asymptomatic infected people in the population), then steady-state prevalence is higher than in the base case; conversely, if the rate is faster, steady-state prevalence is lower. A 10% decrease in this parameter (slower symptom development) increased the optimal capacity by 2, and a 10% increase in this parameter (faster symptom development) decreased the optimal capacity by 2. Changes in transmissibility also had a significant impact on prevalence. For a 10% increase in transmissibility (corresponding to a 10% decrease in the parameter t_1), the optimal capacity increased by 1; for a 10% decrease in transmissibility, the optimal capacity decreased by 2. Changes in the rate at which treated individuals return to the susceptible population ($1/t_3$) and in the rate of exogenous infection ($1/\eta_1$, $1/\eta_2$) had little effect on steady-state prevalence, and thus little effect on the optimal contact tracing capacity.

We performed sensitivity analysis on the network structure by comparing the base case model (which uses a small-world network) to a scale-free network (also known as a preferential-attachment network). We assumed that, although one may not know many details of the contact network, one would likely have a reasonable estimate of disease prevalence in the population, so we used this value as a point of comparison. Thus, we set the average degree of the scale-free network to 2.4 so that the steady-state disease prevalence in the absence of contact tracing would be the same as that for the small-world network with no contact tracing. The degree of 2.4 is less than the value 5 used in the small-world network because, with a scale-free network, the disease spreads more efficiently: the scale-free network has a few individuals with many more contacts than average, and these allow for relatively efficient disease transmission.

Figure 5 shows steady-state annual QALYs experienced for different levels of contact tracing for the small-world network (the same curve as in Figure 3(b)) and for the scale-free network. In both cases, contact tracing has diminishing returns to scale (i.e., the curves are concave). For non-zero levels of contact tracing, disease prevalence is slightly lower in the scale-free network than in the small-world network. This is expected because the scale-free network has some individuals who are highly connected, and finding them (through contact tracing) has a large payoff. For this example, the optimal annual investment is approximately $1000 more for the case of the scale-free network than for the case of the small-world network. This sensitivity analysis highlights the importance of good information about network structure when evaluating how much to invest in contact tracing.

The above sensitivity analyses explore how the optimal budget changes as the network structure and epidemic parameters change. We performed additional sensitivity analyses in which we varied parameter values by up to 200%; multi-way sensitivity analyses in which we varied disease parameters simultaneously; and a stochastic sensitivity analysis in which all parameters were varied within ±10% of their base value. In all cases, contact tracing exhibited diminishing returns to scale as a function of the budget.

4   Contact Tracing and Screening

Thus far, the only form of disease control we have considered is contact tracing. Disease prevalence can also be decreased by screening. One could screen for cases of endogenous infection (cases of

12

infection caused by transmission from individuals in the population) or for cases of exogenous infection (e.g., among immigrants, visitors from other countries, and travelers returning from vacation).

In this section, we address the problem of allocating a combined capacity $K_{total} = K + \lambda$ between contact tracing and screening, and the problem of determining the optimal total capacity K_{total}. Here K is the capacity (manpower) allocated to contact tracing (as in the previous section) and λ is the capacity allocated to screening.

The benefits of contact tracing and screening are larger than the sum of the benefits of doing them separately: the cost effectiveness of contact tracing varies with the amount of screening performed and vice versa. We thus use simulation to determine the optimal mix of contact tracing and screening: we simulate different allocations of a fixed capacity to determine the effectiveness of each combination. Once we know the cost and effectiveness of each combination, we can use the threshold value α to determine the optimal *total* capacity, and the corresponding optimal level of investment in each type of control.

Screening for Exogenous Infection We first consider the case of screening for exogenous infection. Exogenous infection can be a major source of new infection for many diseases: for example, many TB index cases in the US are individuals who have brought the infection from another country. In the US and Canada, long-term immigrants are screened for active TB and HIV as part of the visa process [47, 48].

We assume that with each capacity unit we can either screen 12 people or trace one contact (every $t_4 = t_5 = 5$ days) because the cost of tracing one contact is approximately \$120 (as discussed earlier) and the cost of a gonorrhea test is \$10 [43]. Without any screening, 0.056 exogenous infections occur in the population each day (calculated as $\frac{n}{\eta_1} = \frac{n}{\eta_2} = 500/9000$). We assume that 0.17% of new entrants are infected on average (consistent with gonorrhea prevalence rates in some Asian and eastern European countries [49]); thus, the rate of exogenous infection as a function of λ, the amount of the capacity allocated to screening, is $\eta_1 = \eta_2 = \frac{1}{n}(5/90 - 0.0017\lambda 12/5)$.

Figure 6(a) shows the steady-state prevalence achieved as we vary λ/K_{total} for different total capacities. As one would expect, steady-state prevalence decreases as the total capacity for contact tracing and screening increases. Additionally, we see from this figure that allocating a small fraction of the total capacity to exogenous screening is better for smaller total capacities (no screening,

13

$\lambda = 0$, is optimal for $K_{total} \leq 5$), whereas for larger total capacity it is better to allocate more of the total capacity to exogenous screening.

With this information about the effects of alternative allocations of any fixed total capacity between contact tracing and screening, we can revisit the decision of how large to make the total capacity K_{total}. Figure 6(b) shows the steady-state prevalence as a function of the total capacity, where the capacity is allocated between exogenous screening and tracing so as to minimize the resulting steady-state prevalence (the corresponding minimum from Figure 6(a)). Figure 6(b) also shows the steady-state prevalence as a function of the total capacity, assuming that no screening is used. We see that for capacity $K_{total} > 5$, using a mix of screening and contact tracing achieves lower disease prevalence than does contact tracing alone, and the difference increases as the total capacity increases.

Figure 6(c) shows the total cost of each strategy, including treatment costs (for different levels of K_{total} and different allocations of K_{total} between contact tracing and screening), and the resulting annual steady-state QALYs experienced. The points arranged (more or less) vertically represent different allocations of a given total capacity between screening and contact tracing. The efficient frontier in Figure 6(c) connects the best strategy for each total capacity level. The optimal strategy is given by the point on the curve where the tangent line has a slope of $1/\alpha$. At this point, the annual cost is \$25,500 with approximately 499.1 steady-state QALYs experienced per year. This point corresponds to a capacity K_{total} of 3, all of which is allocated to contact tracing (thus, 3 contacts traced at any one time), with annual disease treatment costs of approximately \$7,500. This is the same solution as was found for the case of contact tracing only.

Screening for Endogenous Infection   We next consider the case of screening for endogenous infection. In our simulation this takes the form of random screening of members of the population. One could think of such screening as resulting from encounters that individuals have with health care providers (either due to symptoms of the disease or for another reason) in which screening is offered. An increase in the endogenous screening rate increases the number of individuals in the population who are screened per unit time, and could correspond to an increase in the rate at which health care providers offer screening to patients. (For example, one practical means of achieving higher rates of routine HIV screening in the US, as recently recommended by the CDC [50], is to

encourage more doctors to routinely offer HIV tests to the patients they see.)

For the case of endogenous screening, we assume that with each capacity unit we can screen 200 people per year. Thus, the mean rate at which individuals move from the infected state (I) to the recovered state (R) when there is endogenous screening at rate λ is $1/t_2 = 1/30 + \lambda 200/n/365$. Figure 7(a) shows the steady-state prevalence achieved as a function of λ/K_{total} for different total capacities K_{total}, and Figure 7(b) shows the steady-state prevalence achieved as a function of the combined capacity K_{total} for endogenous screening and contact tracing. As for the case of exogenous screening, allowing for the possibility of screening for endogenous infection (as occurs in the optimal mix) can reduce prevalence more than contact tracing alone, and the reduction becomes more pronounced as the total capacity increases.

From Figure 7(c) we see that the cost-effectiveness threshold is reached at a point where the total cost is approximately $25,000. This corresponds to a total capacity of $K_{total} = 3$ (annual cost $18,000) plus approximately $7000 in annual treatment costs. From Figure 7(a) we see that for $K_{total} = 3$, approximately one-third of the capacity is allocated to screening and two-thirds is allocated to contact tracing.

5   Example: Contact Tracing for Chlamydia Control

In this section we illustrate the use of our model to evaluate contact tracing for control of chlamydia, a common STD. Estimated chlamydia prevalence in the general US population is about 0.3%, but among young, sexually active individuals (age 15 to 24), prevalence has been found to be 6% or higher [51, 52]. Contact tracing is commonly performed for chlamydia.

We modeled a population of size $n = 500$, reflecting, for example, the size of the sexually active population in a high school. We used the same SIRS model as in Figure 2 (but with different parameter values). We modeled heterosexual transmission of chlamydia, and assumed that the population comprised equal numbers of males and females.

To model this heterosexual population, we created a bipartite graph with equal numbers of males and females. In addition, we modeled low-risk and high-risk males and females, to reflect different levels of risk behavior. To do so, we subdivided the male and female populations into high-contact groups (20% of the total) and low-contact groups (80% of the total).

We assumed an average of three sexual contacts (partnerships) per person. We adjusted this figure upward for high-risk individuals and downward for low-risk individuals. We assumed that the probability that any male-female pair are contacts is independent, and we assumed that these probabilities have a ratio of 7:5:1 for contacts between high-risk individuals, contacts between high-risk and low-risk individuals, and contacts between low-risk individuals, respectively.

We set $t_2 = 15/0.3 = 50$ days to reflect the fact that chlamydia symptoms appear within 1 to 3 weeks after infection, but in 70% of cases, the infection remains asymptomatic [53]. We set the time associated with the sufficient contact rate, t_1, equal to 100. This yields a baseline chlamydia prevalence of 8%, consistent with a study of teenage girls in Philadelphia [51] (before any intervention). Zimmerman et al. [54] found that clients at STD clinics who were found to be infected with chlamydia reported an average of 1.7 contacts. Hence we set the probability that a contact is reported to $1.7/3 = 57\%$. We estimated the cost of treating one case of chlamydia as $c = \$50 + 0.3(\$500) = \$200$: in 70% of cases, an inexpensive ($50) course of antibiotics is sufficient to treat the infection, but in roughly 30% of the cases, patients develop acute pelvic inflammatory disease which must be treated at an additional cost of $500 [55]. For all other parameters of the model (t_3, t_4, t_5, η_1, η_2, and C) we used the same values as shown in Table 1.

Figure 8(a) shows the effect of different contact tracing capacity levels ($K = 0$ to 15) on steady-state chlamydia prevalence in the population. In the absence of any contact tracing, steady-state disease prevalence is 8%. As contact tracing capacity increases, prevalence decreases, but with diminishing returns, as expected. For $K = 5$, steady-state prevalence is 4%; for $K = 10$, steady-state prevalence is approximately 2.7%; for $K = 15$, steady-state prevalence is 2.3%.

Figure 8(b) shows annual costs and QALYs experienced for each level of contact tracing. For a cost-effectiveness threshold of $\alpha = \$50,000/QALY$ gained, the optimal contact tracing capacity is $K = 11$. The annual cost is approximately $92,000, corresponding to $66,000 for contact tracing plus $26,000 for treatment. The resulting endemic level of disease prevalence is approximately 2.6% (Figure 8(a)). If the public health budget for this example constrains contact tracing to fewer than 11 individuals being traced simultaneously, Figure 8(b) shows that significant health benefits can still be achieved by contact tracing. For example, a contact tracing capacity of $K = 6$ will still reduce steady-state prevalence from 8% to 3.7%. A simulation model such as ours allows one to quantify the tradeoffs associated with different levels of contact tracing.

6   Discussion

We have presented a general framework for evaluating the optimal level of investment in contact tracing. This framework combines concepts from cost-effectiveness analysis with simulation of a disease and the effects of contact tracing among a network of individuals. Our simulation results suggest that contact tracing is likely to have diminishing returns to scale: incremental increases in the budget for contact tracing yield diminishing decreases in the disease prevalence. This makes intuitive sense: the larger the number of contacts traced per unit time, the less likely it is that the incremental contacts traced will be infected, and thus the smaller the number of new cases that will be identified and removed. Use of a cost-effectiveness framework, combined with the simulation results, allows one to determine the appropriate level of investment in contact tracing. We also showed how, when other interventions such as screening are available, simulation can be used to determine the best mix of interventions for any given capacity level, and then cost-effectiveness principles can be applied to determine the appropriate total capacity for screening and contact tracing.

Our results are based on a limited set of simulations. Further analyses could explore the robustness of our findings under different conditions: for example, for different networks, diseases, epidemic models, disease parameters, and populations.

One useful avenue for future research is to extend our simulation model to capture more details of contact tracing and disease transmission and progression. For example, our current model stylizes the screening of infections from exogenous sources. A more realistic model could break out the various sources of exogenous infection (e.g., holiday travelers, visitors from certain countries, legal immigrants, and illegal immigrants) and the opportunities to screen them (e.g., when they request a visa or at clinics in immigrant neighborhoods). As another example, our model of contact tracing does not include the genotype information available to investigators that allows them to distinguish between new and continuing outbreaks. In practice, when a new outbreak of a disease is detected, the intensity of contact tracing is often increased until a significant level of epidemic control has been achieved. Thus, a natural extension of our work is to consider the case of dynamically changing levels of contact tracing.

A more sophisticated model would show additional benefits of contact tracing not captured by

our model. As contact tracing capacity increases, the average time from acquisition of infection to detection and treatment decreases. Our model captures this effect (Figure 9). However, our model does not capture the fact that the most severe symptoms and complications are in patients whose disease was untreated the longest. Extending our model to include multiple infected states (for example, states representing benign versus acute symptoms), which differ in their infectiousness, quality multipliers, and treatment costs, would model this phenomenon and thus help quantify this additional benefit of contact tracing.

The simplicity of our current model also means that the selection of contacts to trace is extremely crude and does not reflect all the considerations (such as stage of infection, strength of immune system, age, demographics, etc.) that are used in the real world. By assuming that contacts are located after a constant number of days, our model ignores the possibility of not locating a contact, the uncertainty in the time required to locate a contact, and the option of giving up the search. A model incorporating such features would be more realistic and could be used to compare different strategies for selecting who to trace.

Another useful avenue for further research would be to tailor the analysis to specific diseases of interest. A tailored model could be used to determine the appropriate level of contact tracing for a specific disease in a specific region. The simulation model we have presented here could provide the foundation for such a policy tool. As an example, we applied our SIRS model to show how one could examine different levels of contact tracing for chlamydia control. For other diseases, a different epidemic model (with appropriate adjustments to the network model) might be needed.

For example, for TB, the disease model should include latent and active infection stages, with disease progression times set appropriately. Further, the contact network should allow contacts of greater and lesser strength (e.g., family members versus coworkers in a well-ventilated office). In our simulation, a contact's priority score is an indicator of the likelihood that this contact is infected. To better model TB contact tracing, it would be useful to distinguish individuals by their potential chance of acquiring infection, as is done in practice. For example, TB contact tracing in the US gives priority to contacts who are children or who have AIDS.

To model HIV and some STDs, the disease model should include asymptomatic and symptomatic disease stages. For HIV, it might be appropriate to incorporate several modes of transmission (e.g., heterosexual partnerships, same-sex partnerships, and needle-sharing partnerships). In

addition, use of a dynamic contact network would reflect the pair formation and dissolution that occurs in social networks of such diseases (see, for example, Kretzschmar [21]).

More work is also needed to understand the network structure underlying a particular epidemic in a particular population. Some limited work has been done to characterize specific networks of infectious disease (e.g., Klovdahl et al. [34], Weeks et al. [35], Parker et al. [36], Helleringer et al. [37]). Further knowledge of network structure, and related disease transmission features such as pair formation and dissolution, is crucial to evaluating the effectiveness, and cost-effectiveness, of different investments in contact tracing in any given setting.

We have considered contact tracing in an endemic disease setting. Contact tracing is also important (perhaps more so) for containing outbreaks of epidemic diseases. Then the crucial question is whether one can find the contacts faster than they can spread the disease, or win the "race to trace" as Kaplan et al. [56] put it. For this problem, the decision maker's goal is likely to be that of determining the minimum level of contact tracing that is needed to reduce the reproductive rate of infection below 1, perhaps within a specified period of time. This approach may be relevant for particularly virulent diseases such as smallpox [56], extremely drug-resistant TB [57], or a gonorrhea "superbug" [58].

Contact tracing can be an effective means of disease control, but it is only useful up to a point because incremental increases in the level of contact tracing are likely to yield diminishing benefits. Simulation can be used to estimate the benefits of contact tracing as a function of its intensity. Then, cost-effectiveness analysis can be used to determine the optimal level of investment in contact tracing (or the optimal level of investment in contact tracing and screening). Such analysis can help public health departments make the most cost-effective use of their available funds for disease control.

References

[1] Saretok P, Brouwers L (2007) Microsimulation of pandemic influenza in Sweden. In: Proceedings of the 2007 Western Multiconference on Computer Simulation, pp 73–78

[2] Bozzette SA, Boer R, Bhatnagar V, Brower JL, Keeler EB, Morton SC, Stoto MA (2003) A model for a smallpox-vaccination policy. N Engl J Med 348(5):416–425

[3] Hethcote HW (1982) Gonorrhea modeling: a comparison of control methods. Math Biosci 58: 93–109

[4] Zaric GS, Brandeau ML (2001) Resource allocation for epidemic control over short time horizons. Math Biosci 171:33–58

[5] Zaric GS, Brandeau ML (2002) Dynamic resource allocation for epidemic control in multiple populations. IMA J Math Appl Med Bio 19:235–255

[6] Longini IM Jr, Ackerman E, Elveback L (1978) An optimization model for influenza A epidemics. Math Biosci 38:141–157

[7] Halloran ME, Longini IM Jr, Nizam A, Yang Y (2002) Containing bioterrorist smallpox. Science 298(5597):1428–1432

[8] Müller J (1998) Optimal vaccination patterns in age-structured populations. SIAM J Appl Math 59:222–241

[9] Revelle C, Feldmann F, Lynn W (1969) An optimization model tuberculosis epidemiology. Manage Sci 16(4):B190–B211

[10] Centers for Disease Control and Prevention (2000) Chapter 10 - community TB control: Identification of persons who have clinically active TB, 2000. URL http://www.cdc.gov/nchstp/tb/pubs/corecurr/Chapter10/Chapter_10_Identification.htm

[11] Centers for Disease Control and Prevention (2002) Sexually transmitted diseases treatment guidelines 2002. URL http://www.cdc.gov/STD/treatment/1-2002TG.htm

[12] Cowan FM, French R, Johnson AM (1996) The role and effectiveness of partner notification in STD control: a review. Genitourin Med 72:247–252

[13] Clarke J (1998) Contact tracing for chlamydia: data on effectiveness. Int J STD AIDS 9: 187–191

[14] Lipsitch M, Cohen T, Cooper B, Robins JM, Ma S, James L, Gopalakrishna G, Chew SK, Tan CC, Samore MH, Fisman D, Murray M (2003) Transmission dynamics and control of severe acute respiratory syndrome. Science 300(5627):1966–1970

[15] Kiss IZ, Green DM, Kao RR (2005) Disease contact tracing in random and clustered networks. Proc R Soc Lond B Biol Sci 272:1407–1414

[16] Porco TC, Holbrook KA, Fernyak SE, Portnoy DL, Reiter R, Aragón TJ (2004) Logistics of community smallpox control through contact tracing and ring vaccination: a stochastic network model. BMC Public Health 4(34)

[17] Kretzschmar M, van den Hof S, Wallinga J, van Wijngaarden J (2004) Ring vaccination and smallpox control. Emerg Infect Dis 10(5)

[18] Wu JT, Riley S, Fraser C, Leung GM (2006) Reducing the impact of the next influenza pandemic using household-based public health interventions. PLoS Medicine 3:1532–1540

[19] Hyman JM, Li J, Stanley EA (2003) Modeling the impact of random screening and contact tracing in reducing the spread of HIV. Math Biosci 181:17–54

[20] Armbruster B, Brandeau ML (2007) Optimal mix of screening and contact tracing for endemic diseases. Math Biosci 209:386–402

[21] Kretzschmar M (2000) Sexual network structure and sexually transmitted disease prevention: a modeling perspective. Sex Transm Dis 27:627–635

[22] Müller J, Kretzschmar M, Dietz K (2000) Contact tracing in stochastic and deterministic epidemic models. Math Biosci 164:39–64

[23] Huerta R, Tsimring LS (2002) Contact tracing and epidemics control in social networks. Phys Rev E Stat Nonlin Soft Matter Phys 66(5):056115

[24] Eames KTD, Keeling MJ (2003) Contact tracing and disease control. Proc R Soc Lond B Biol Sci 270:2565–2571

[25] Dasgupta K, Schwartzman K, Marchand R, Tennenbaum TN, Brassard P, Menzies D (2000) Comparison of cost-effectiveness of tuberculosis screening of close contacts and foreign-born populations. Am J Respir Crit Care Med 162(6):2079–2086

[26] Macintyre CR, Plant AJ, Hendrie D (2000) The cost-effectiveness of evidence-based guidelines and practice for screening and prevention of tuberculosis. Health Econ 9(5):411–421

[27] Varghese B, Peterman TA, Holtgrave DR (1999) Cost-effectiveness of counseling and testing and partner notification: a decision analysis. AIDS 13(13):1745–1751

[28] Cohen DA, Wu SY, Farley TA (2004) Comparing the cost-effectiveness of HIV prevention interventions. J Acquir Immune Defic Syndr 37(3):1404–1414

[29] Howell MR, Kassler WJ, Haddix A (1997) Partner notification to prevent pelvic inflammatory disease in women: cost-effectiveness of two strategies. Sex Transm Dis 24(5):287–292

[30] Oxman GL, Doyle L (1996) A comparison of the case-finding effectiveness and average costs of screening and partner notification. Sex Transm Dis 23(1):51–57

[31] Welte R, Kretzschmar M, Postma MJ, Leidl R, van den Hoek JAR, Jager JC, Postma MJ (2000) Cost-effectiveness of screening programs for Chlamydia trachomatis: a population based dynamic approach. Sex Transm Dis 27:518–529

[32] Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. Nature 393: 440–442

[33] Schneeberger A, Mercer CH, Gregson SA, Ferguson NM, Nyamukapa CA, Anderson RM, Johnson AM, Garnett GP (2004) Scale-free networks and sexually transmitted diseases: a description of observed patterns of sexual contacts in Britain and Zimbabwe. Sex Transm Dis 31(6):380–387

[34] Klovdahl AS, Potterat JJ, Woodhouse DE, Muth JB, Muth SQ, Darrow WW (1994) Social networks and infectious disease: the Colorado Springs study. Soc Sci Med 38(1):79–88

[35] Weeks MR, Clair S, Borgatti SP, Radda K, Schensul JJ (2002) Social networks of drug users in high-risk sites: finding the connections. AIDS Behav 6(2):194–206

[36] Parker M, Ward H, Day S (1998) Sexual networks and the transmission of HIV in London. J Biosoc Sci 30:63–83

[37] Helleringer S, Kohler HP, Chimbiri A, Chatonda P, Mkandawire J (2007) The Likoma network study: context, data collection and survey instruments. URL http://www.ssc.upenn.edu/~hpkohler/working-papers/hell06likdoc.pdf. working paper

[38] Wylie JL, Jolly A (2001) Patterns of chlamydia and gonorrhea infection in sexual networks in Manitoba, Canada. Sex Transm Dis 28(1):14–24

[39] Armbruster B, Brandeau ML (2007) Who do you know? A simulation study of infectious disease control through contact tracing. In: Proceedings of the 2007 Western Multiconference on Computer Simulation, pp 79–85

[40] Gold M, Siegel J, Russell L, Weinstein M (1996) Cost-effectiveness in health and medicine, Oxford University Press

[41] National Institutes of Health (2001) Workshop summary: Scientific evidence on condom effectiveness for sexually transmitted disease (STD) prevention. URL http://www.niaid.nih.gov/dmid/stds/condomreport.pdf.

[42] Centers for Disease Control and Prevention (2006) Fact sheet: Gonorrhea. URL http://www.cdc.gov/std/Gonorrhea/STDFact-gonorrhea.htm.

[43] Mehta SD, Bishai D, Howell MR, Rothman RE, Quinn TC, Zenilman JM (2002) Cost-effectiveness of five strategies for gonorrhea and chlamydia control among female and male emergency department patients. Sex Transm Dis 29(2):83–91

[44] Centers for Disease Control and Prevention (2001) Program Evaluation. URL http://www.cdc.gov/std/program/training/TOC-PGtraining.htm.

[45] Neumann P et al. (2001) Preference weights 1998–2001. URL http://www.tufts-nemc.org/cearegistry/data/docs/phaseIIpreferenceweights.pdf.

[46] Owens DK (1998) Interpretation of cost-effectiveness analyses. J Gen Intern Med 13:716–717

[47] U.S. Citizenship and Immigration Services (2006) Medical examinations. URL http://www.uscis.gov/propub/ProPubVAP.jsp?dockey=1c56d78b831659735c9fd1c1f7dbd68f

[48] Citizenship and Immigration Canada (2002) Fact sheet — medical testing and surveillance. URL http://www.cic.gc.ca/english/pub/fs-medical.html

[49] World Health Organization (2001) Global prevalence and incidence of selected curable sexually transmitted infections: Overview and estimates. URL http://www.who.int/docstore/hiv/GRSTI/index.htm

[50] Centers for Disease Control and Prevention (2006) Revised recommendations for HIV testing of adults, adolescents, and pregnant women in health-care settings. MMWR 55(RR14):1–17

[51] Webb AL (2004) But I didn't know..., 2004. Philadelphia City Paper January 22

[52] Centers for Disease Control and Prevention (2005) STD surveillance 2005 national profile – chlamydia. URL http://www.cdc.gov/std/stats/chlamydia.htm.

[53] National Institutes of Health (2006) Chlamydia. URL http://www.niaid.nih.gov/factsheets/stdclam.htm.

[54] Zimmerman HL, Potterat JJ, Dukes RL, Muth JB, Zimmerman HP, Fogle JS, Pratts CI (1990) Epidemiologic differences between chlamydia and gonorrhea. Am J Public Health, 80 (11):1338–1342

[55] Hu D, Hook EW 3rd, Goldie SJ (2004) Screening for Chlamydia trachomatis in women 15 to 29 years of age: a cost-effectiveness analysis. Ann Intern Med 141:501–513

[56] Kaplan EH, Craft DL, Wein LM (2002) Emergency response to a smallpox attack: the case for mass vaccination. PNAS 99(16):10935–10940

[57] Centers for Disease Control and Prevention (2007) Extensively drug-resistant tuberculosis — United States, 1993–2006. MMWR 56(11):250–253

[58] Centers for Disease Control and Prevention (2007) Update to CDC's sexually transmitted diseases treatment guidelines, 2006: fluoroquinolones no longer recommended for treatment of gonococcal infections. MMWR 56(14):332–336
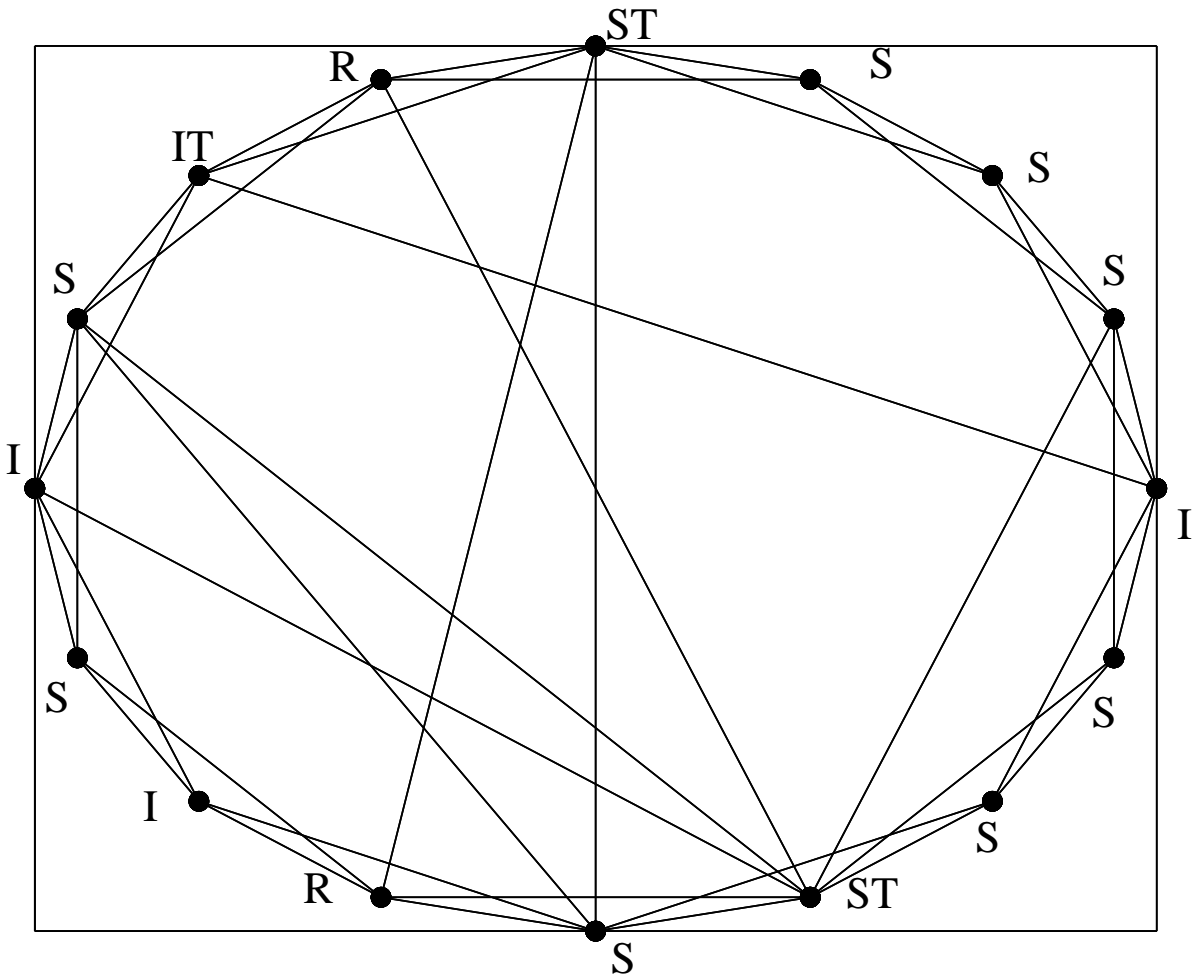
Figure 1: A small-world graph with nodes in various states. The states are: susceptible (S), infected (I), removed (R), susceptible and being traced (ST), and infected and being traced (IT).
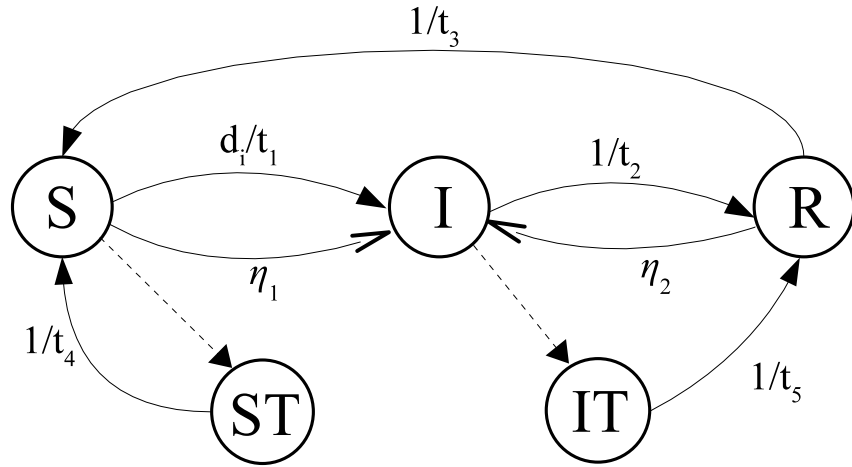
Figure 2: Possible states of an individual and the transition times between them. The states are: susceptible (S), infected (I), removed (R), susceptible and being traced (ST), and infected and being traced (IT). The dashed arrows mark the instantaneous transitions $S \to ST$ and $I \to IT$ that occur when we decide to trace this individual. The terms $t_0, \ldots, t_5$ are time constants, d_i is the number of infected neighbors of this individual, and η_1 and η_2 are rates of exogenous infection.
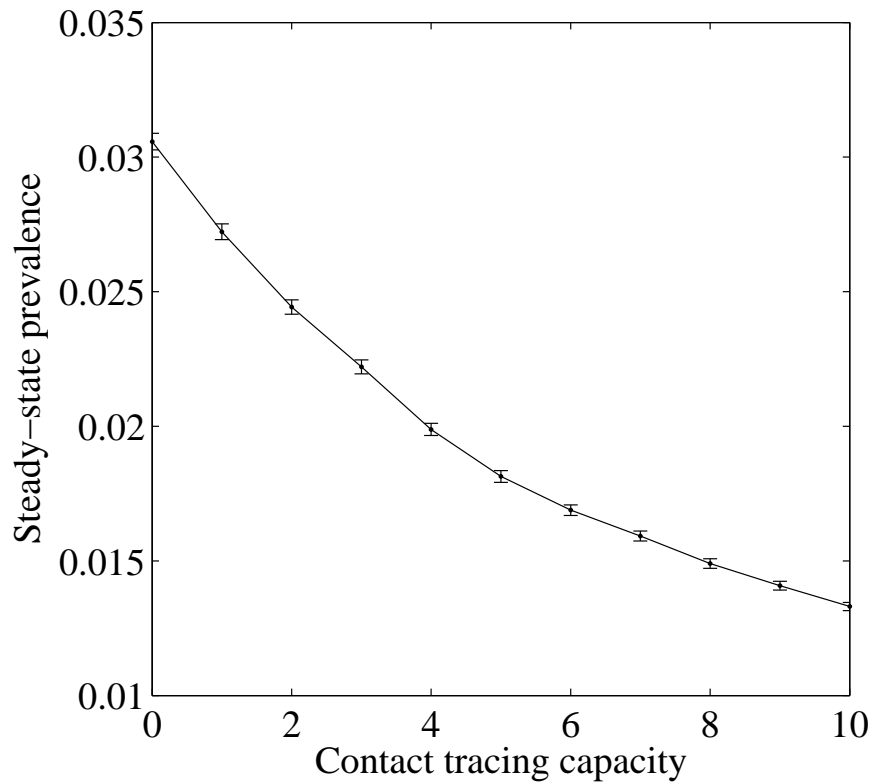
(a) Simulated steady-state disease prevalence as a function of contact tracing capacity, K.
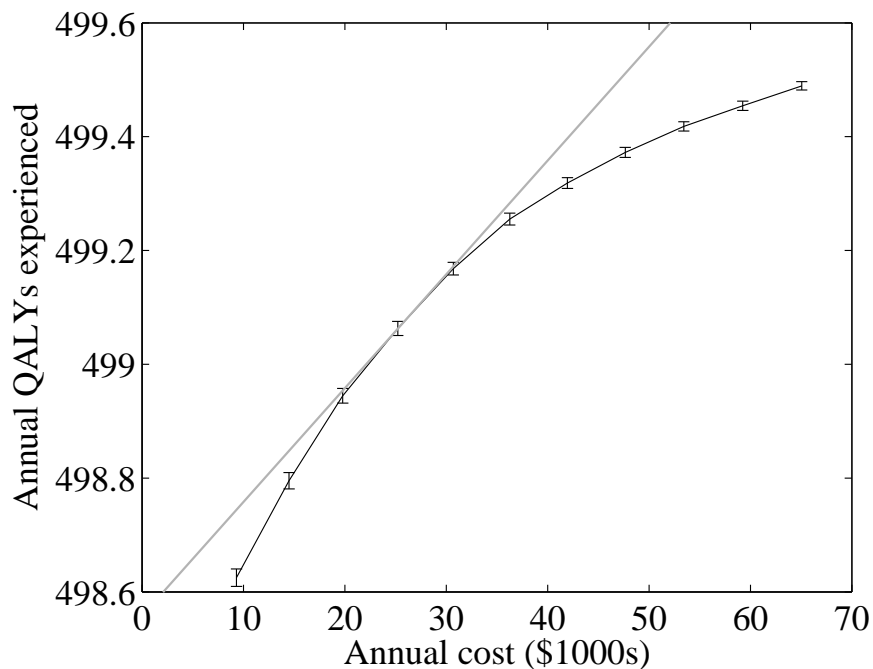
(b) Steady-state annual cost and QALYs experienced as a function of contact tracing capacity. The points correspond to capacity K ranging from 0 to 10, as in Figure 3(a). At the optimal budget this curve has a slope of $1/\alpha$ (the slope of the gray line). Here $\alpha = \$50,000/\text{QALY}$.

Figure 3: Effect of varying the contact tracing capacity, K. Vertical bars represent 95% confidence intervals (1600 runs).
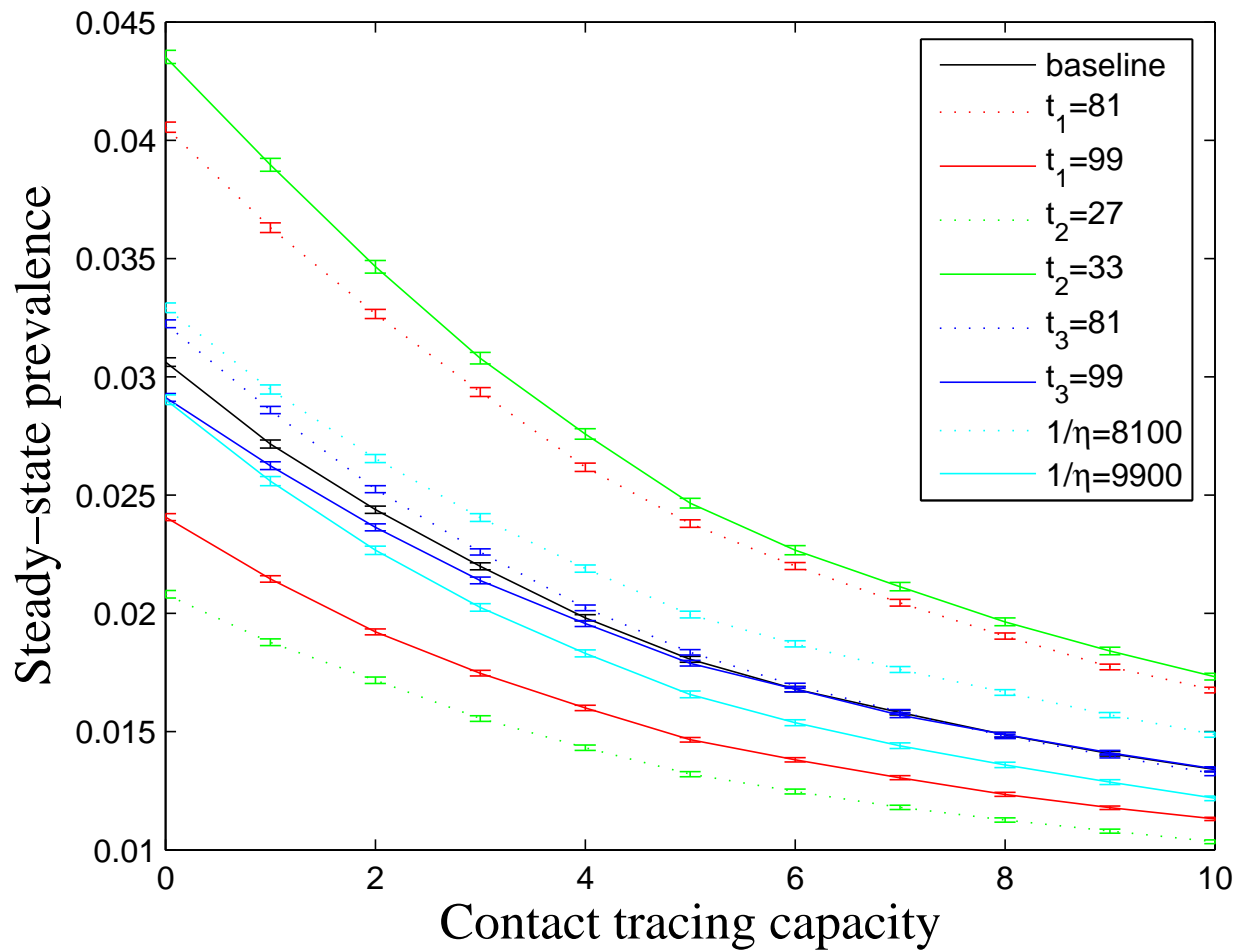
Figure 4: 1-way sensitivity analysis: steady-state disease prevalence as a function of contact tracing capacity, K. Vertical bars represent 95% confidence intervals (1600 runs).
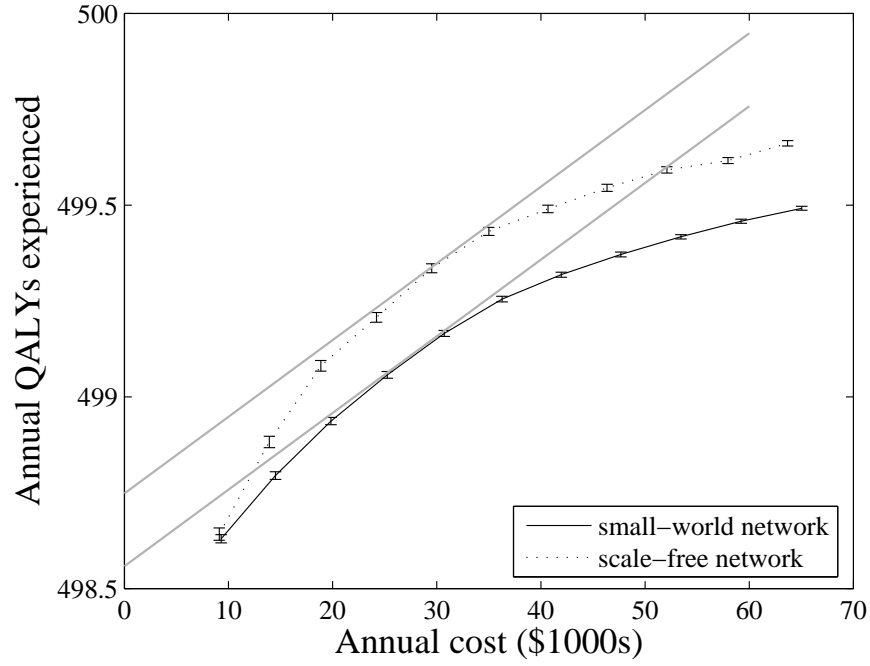
Figure 5: Comparison of the cost-effectiveness of contact tracing on a small-world network (as in Figure 3(b)) and on a scale-free network of the same size ($n = 500$) whose average degree (2.4) is chosen so that the steady-state prevalence without contact tracing is similar to that for the small-world network. The points correspond to the contact tracing capacity K ranging from 0 to 10. At the optimal budget this curve has a slope of $1/\alpha$ (the slope of the gray lines). Here $\alpha = \$50,000/\text{QALY}$. Vertical bars represent 95% confidence intervals (1200 runs).

(a) Steady-state prevalence as a function of the fraction spent on screening λ/K_{total} for various values of the total capacity K_{total}.

(b) Steady-state prevalence achieved as a function of the total capacity K_{total} for screening and contact tracing. The solid line allocates the capacity optimally while the dotted line is from figure 3(a) where we used no screening ($\lambda = 0$).
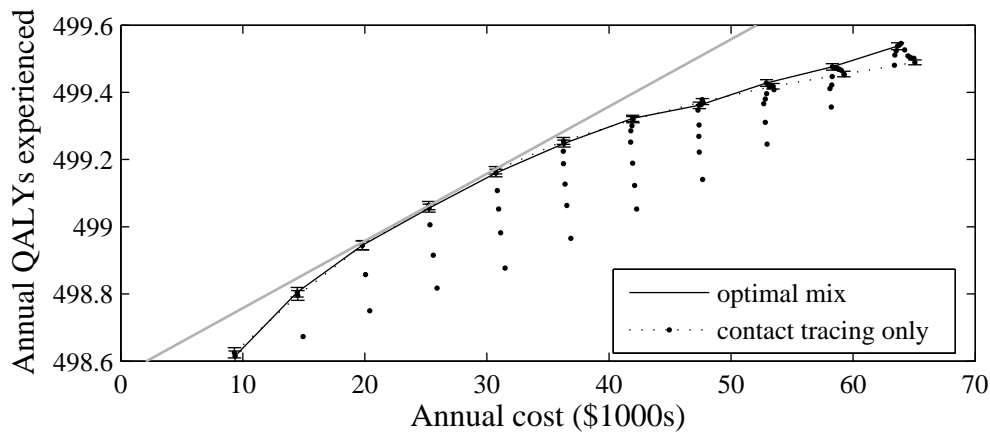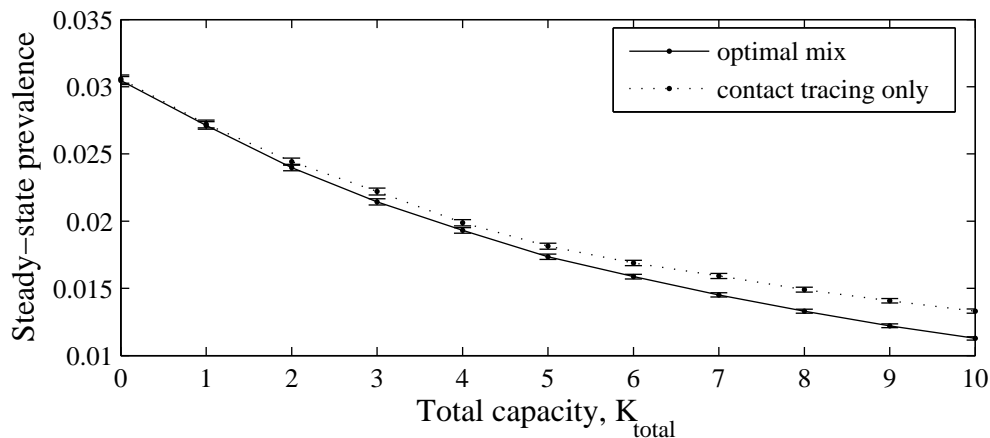

(c) Steady-state annual cost and QALYs experienced as a function of the total capacity K_{total} for screening and contact tracing. The points arranged (more or less) vertically represent different allocations of a given total capacity, K_{total}, between screening and contact tracing. At the optimal budget, this curve has slope equal to $1/\alpha$ (the slope of the gray line). Here $\alpha = \$50,000/QALY$.
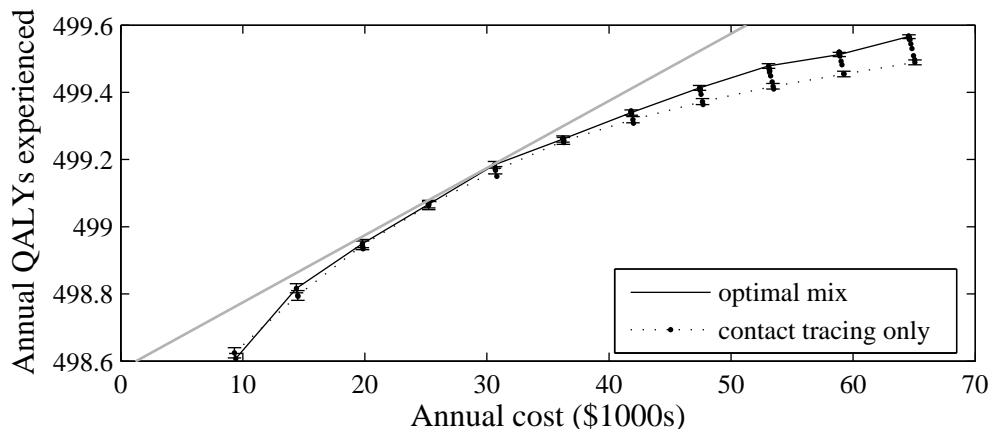
Figure 6: Effects of allocating a total capacity $K_{total} = K + \lambda$ between contact tracing, K, and screening for *exogenous* infections, λ. Vertical bars represent 95% confidence intervals (1600 runs).

(a) Steady-state prevalence as a function of the fraction spent on screening λ/K_{total} for various values of the total capacity K_{total}.
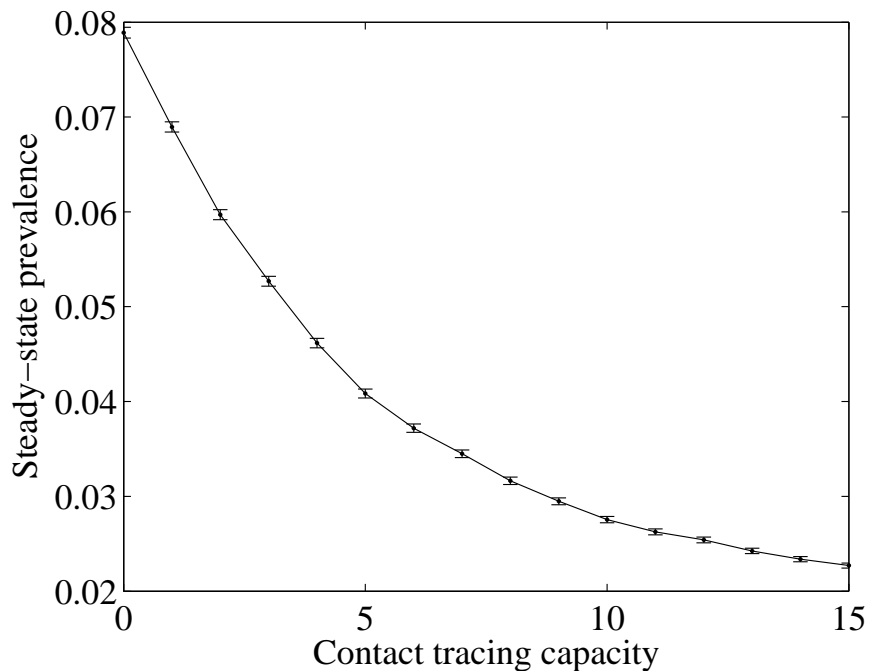

(b) Steady-state prevalence achieved as a function of the total capacity K_{total} for screening and contact tracing. The solid line allocates the capacity optimally while the dotted line is from figure 3(a) where we used no screening ($\lambda = 0$).
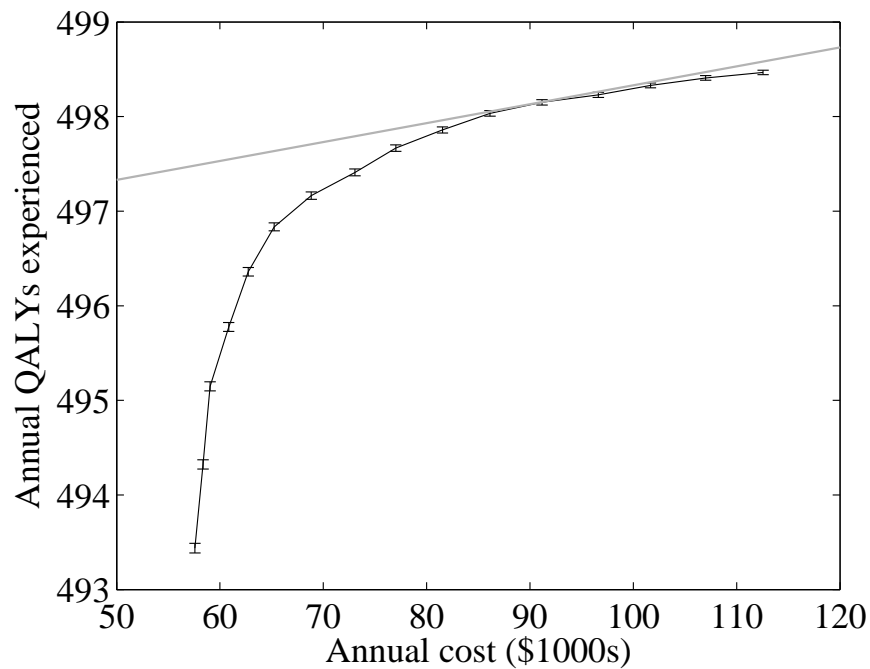

(c) Steady-state annual cost and QALYs experienced as a function of the total capacity K_{total} for screening and contact tracing. The points arranged (more or less) vertically represent different allocations of a given total capacity, K_{total}, between screening and contact tracing. At the optimal budget, this curve has slope equal to $1/\alpha$ (the slope of the gray line). Here $\alpha = \$50,000/QALY$.

Figure 7: Effects of allocating a total capacity $K_{total} = K + \lambda$ between contact tracing, K, and screening for *endogenous* infections, λ. Vertical bars represent 95% confidence intervals (1600 runs).

(a) Simulated steady-state disease prevalence as a function of contact tracing capacity, K, for the case of chlamydia.

(b) Steady-state annual cost and QALYs experienced as a function of contact tracing capacity. The points correspond to capacity K ranging from 0 to 15, as in Figure 8(a). At the optimal budget, this curve has a slope of $1/\alpha$ (the slope of the gray line). Here $\alpha = \$50,000/QALY$.

Figure 8: Evaluating contact tracing capacity for chlamydia control. Vertical bars represent 95% confidence intervals (1600 runs).
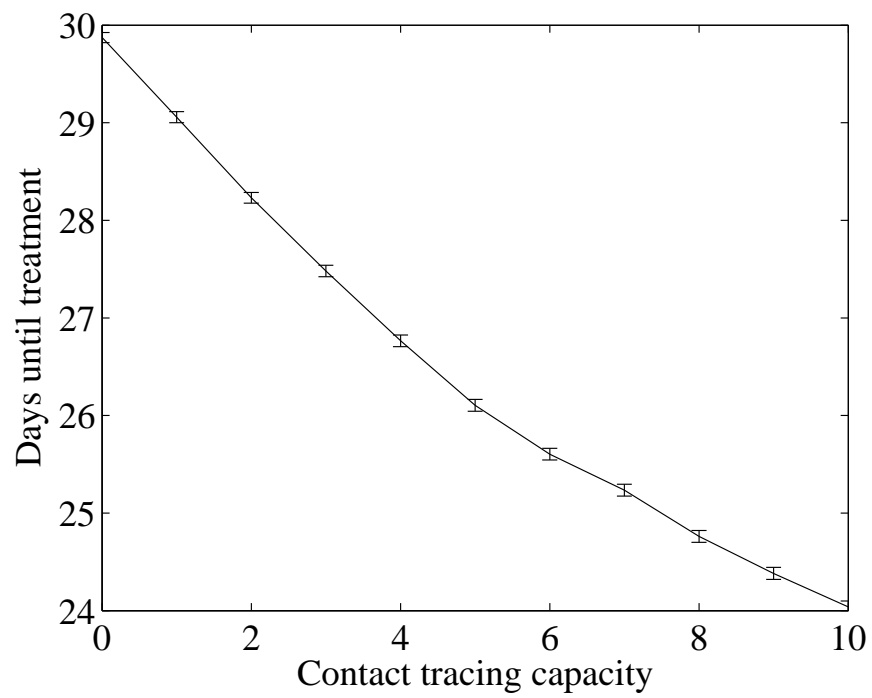
Figure 9: Average time from acquisition of infection until treatment is initiated, as a function of contact tracing capacity, K. Vertical bars represent 95% confidence intervals (1600 runs).