Mathematical Approaches to Infectious Disease Prediction and Control

Nedialko B. Dimitrov

Operations Research Department, Naval Postgraduate School, Monterey, California
ned@alumni.cs.utexas.edu

Lauren Ancel Meyers

Section of Integrative Biology, The University of Texas at Austin, Austin, Texas
Santa Fe Institute, Santa Fe, New Mexico
laurenmeyers@mail.utexas.edu

Abstract

Mathematics has long been an important tool for understanding and controlling the spread of infectious diseases. Here, we begin with an overview of compartmental models, the traditional approach to modeling infectious disease dynamics, and then introduce contact network epidemiology, a relatively new approach that applies bond percolation on random graphs to model the spread of infectious disease through heterogeneous populations. As we illustrate, these methods can be used to address public health challenges and have recently been coupled with powerful computational methods to optimize epidemic control strategies.

1 Introduction

As the novel 2009 H1N1 strain of influenza emerged out of Mexico and rapidly spread around the globe, public health agencies scrambled to understand and control its spread. The World Health Organization (WHO) and the US Centers for Disease Control and Prevention (CDC) immediately looked to statisticians and mathematical modelers to make sense of the sparse and noisy data available from the initial outbreaks. In order to make informed decisions about school closures, travel restrictions, and uses of limited resources such as antiviral medications, public health authorities needed to know the rate of spread and severity of the new strain, as well as whether prior flu vaccines or exposure to existing strains provided any immunity to H1N1/09. As evidenced by several papers published jointly by academic researchers and public health officials within a few months of the emergence of H1N1/09 [23, 47, 65], mathematical modeling played a vital role in shaping the global public health response to the pandemic.

   The mathematical techniques used to understand, forecast and control the spread of infectious diseases like influenza are diverse and growing rapidly. Some techniques have been newly developed while others build upon existing methods from diverse fields including dynamical systems, stochastic processes, statistical physics, graph theory, statistics, operations research and high performance computing. Here, we present an overview of the some of the most widely used and promising mathematical approaches to modeling the spread of infectious disease. In Section 2, we present an overview of compartmental models, the workhorse of mathematical epidemiology throughout the twentieth century. In Sections 3 and 4, we discuss parameterization of infectious disease models and some limitations of the standard modeling approaches. In Section 5, we turn to contact network modeling, a relatively new analytical approach that explicitly considers complex human population structures, thereby overcoming a major limitation of compartmental models. In Section 6, we conclude by discussing the promising application of advanced optimization methods to designing public health policies.

2 Compartmental SIR Model

The first differential equation models of infectious disease dynamics go back as far as 1766, to the work of Daniel Bernoulli, which has been recently re-published [12]. Modern differential equation models of epidemics were introduced by Kermack and McKendrick [30] and later expanded by Anderson and May [4, 5].

   In this section, we present an intuitive overview of modern compartmental models. First we discuss the relationship between discrete-time, discrete-state population models and the continuous-time, continuous-state compartmental models; second, we examine the critical parameters required to instantiate epidemic models. Throughout this section, we focus on a simple and widely-used version of the SIR compartmental model and provide insights into the model's behavior.

   Consider a population of N individuals and the following simple discrete-time, discrete-state epidemic model. Each individual begins in one of the three possible states:

1. *Susceptible*, meaning that the individual has never had the disease and is susceptible to being infected.
2. *Infected*, meaning that the individual currently has the disease and can infect other people.
3. *Resistant*, meaning that the individual does not have the disease, cannot infect others, and cannot be infected.

The model then evolves in discrete time steps, with all individuals simultaneously acting as follows in each time step:

1. Each susceptible individual draws a uniformly random person from the population. If the person drawn is infected, then the susceptible individual changes his state to infected with probability β.
2. Each infected individual changes his state to resistant with probability γ.
3. Each resistant individual remains resistant.

Intuitively, the above discrete-time, discrete-space model simulates a population of interacting individuals. Interactions are modeled from the perspective of susceptible individuals, who can become infected during interactions with infected individuals. A population that interacts in such a uniformly random and independent way between time steps is called a *homogeneously mixed population*.

The model also simulates the progression of the disease through the three available states. Individuals are first susceptible, then infected, and then become resistant by acquiring immunity to the disease. The parameter β captures the ability of the disease to be transmitted from one person to another, while the parameter γ is related to length of the period for which an individual can transmit the disease, called the *infectious period*. Specifically, the total time spent in the infected state by an individual is a geometric random variable with success probability γ, making the expected length of the infectious period equal to $\frac{1}{\gamma}$.

The abbreviation SIR stands for the three available states: susceptible, infected, and resistant. However, the term *SIR model* typically refers to a continuous differential equation model that we will now derive from the above discrete model. Suppose that the initial condition of the population is given, and let the random variables $X(t)$, $Y(t)$ and $Z(t)$ denote the number of susceptible, infected, and resistant individuals in the population at time t. Since each individual is always in one of the three states, it is always the case that $X(t) + Y(t) + Z(t) = N$. Given the values of these three quantities at time t, we can calculate their expected values at time $t + 1$:

$$E[X(t+1)] = X(t) \cdot (1 - \frac{Y(t)}{N} \cdot \beta)$$

$$E[Y(t+1)] = X(t) \cdot (\frac{Y(t)}{N} \cdot \beta) + Y(t) \cdot (1 - \gamma)$$

$$E[Z(t+1)] = Y(t) \cdot \gamma + Z(t).$$

These equations are based on the basic assumptions of the model. The first equation expresses that $X(t)$ susceptible individuals at time t act similarly and independently, each remaining in the susceptible state with probability $(1 - \frac{Y(t)}{N} \cdot \beta)$; the second equation indicates that an expected $X(t) \cdot (\frac{Y(t)}{N} \cdot \beta)$ individuals enter the infected state and $Y(t) \cdot \gamma$ infected individuals leave the infected state; and the third equation reflects that the individuals who leave the infected state enter the resistant state, and resistant individuals remain resistant. Re-arranging the three equations above, we obtain

$$E[X(t+1)] - X(t) = -\beta \cdot X(t) \cdot \frac{Y(t)}{N}$$

$$E[Y(t+1)] - Y(t) = \beta \cdot X(t) \cdot \frac{Y(t)}{N} - \gamma \cdot Y(t)$$

$$E[Z(t+1)] - Z(t) = \gamma \cdot Y(t).$$

The next step of the derivation uses the mean-field approximation, which is at times mathematically controversial. The controversy stems from the fact that the accuracy of the approximation is application dependent and often difficult to analyze. The *mean-field approximation* allows us to forget the fact that $X(t), Y(t)$, and $Z(t)$ are random variables and simply equate them with their

expectations. Regardless of the controversy, the mean-field approximation is practically useful and later in this section we provide some experimental evidence of its applicability to our homogeneously mixed population model. Applying the mean-field approximation to the above equations, allows us to rewrite them as:

$$X(t+1) - X(t) = -\beta \cdot X(t) \cdot \frac{Y(t)}{N}$$

$$Y(t+1) - Y(t) = \beta \cdot X(t) \cdot \frac{Y(t)}{N} - \gamma \cdot Y(t)$$

$$Z(t+1) - Z(t) = \gamma \cdot Y(t).$$

The differential equations of the continuous-time, continuous-state SIR model that we seek are evident in the three difference equations above. However, the above equations are fixed at a time-difference of 1, because our discrete-time model moves in these increments. To complete the final step, we create a sequence of discrete-time models from which we derive the continuous-time model. The derivation is reminiscent of the proof that a sequence of geometric random variables converges to an exponential random variable (see Appendix A).

Let Δt be a real number less than one. We create a discrete-time model that moves in increments of Δt simply by changing the parameters β, γ to $\Delta t\beta, \Delta t\gamma$. Intuitively, by thus altering the parameters, we keep the expected number of successes in a unit time interval the same for each discrete-time model. Applying our derived difference equations above, we have

$$X(t+\Delta t) - X(t) = -\Delta t\beta \cdot X(t) \cdot \frac{Y(t)}{N}$$

$$Y(t+\Delta t) - Y(t) = \Delta t\beta \cdot X(t) \cdot \frac{Y(t)}{N} - \Delta t\gamma \cdot Y(t)$$

$$Z(t+\Delta t) - Z(t) = \Delta t\gamma \cdot Y(t).$$

Dividing by Δt and taking the limit as Δt goes to zero gives us the following differential equations:

$$\frac{dX(t)}{dt} = -\beta \cdot X(t) \cdot \frac{Y(t)}{N} \tag{1}$$

$$\frac{dY(t)}{dt} = \beta \cdot X(t) \cdot \frac{Y(t)}{N} - \gamma \cdot Y(t) \cdot \tag{2}$$

$$\frac{dZ(t)}{dt} = \gamma \cdot Y(t). \tag{3}$$

Equations (1)-(3) are referred to as a *mass-action SIR compartmental model*. The adjective mass-action comes from the fact that in the original discrete model, all individuals act similarly yet separately from each other. The adjective compartmental comes from viewing the three disease states as compartments into and out of which individuals move throughout the epidemic. Finally, if we make the variable substitutions $S(t) = \frac{X(t)}{N}, I(t) = \frac{Y(t)}{N}$ and $R(t) = \frac{Z(t)}{N}$, we have

$$\frac{dS(t)}{dt} = -\beta \cdot S(t) \cdot I(t) \tag{4}$$

$$\frac{dI(t)}{dt} = \beta \cdot S(t) \cdot I(t) - \gamma \cdot I(t) \tag{5}$$
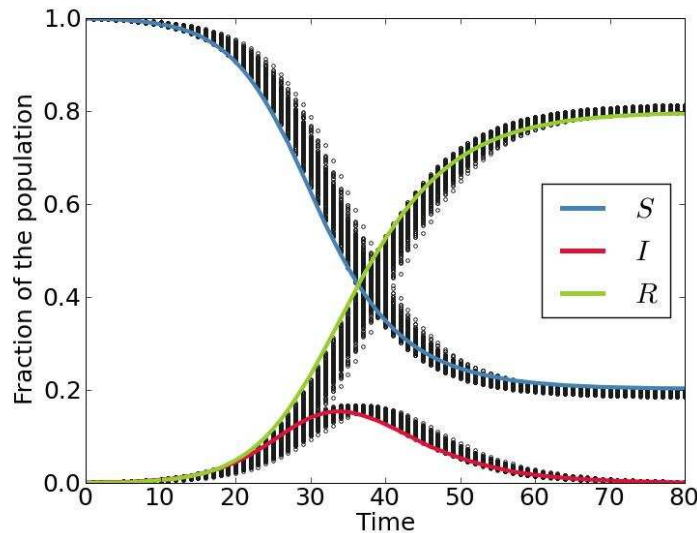
$$\frac{dR(t)}{dt} = \gamma \cdot I(t). \tag{6}$$

Figure 1: Comparison of a discrete-time, discrete-state disease model and the corresponding compartmental SIR model. Both models have parameters $\beta, \gamma = 0.4, 0.2$. The discrete model has population size of $100,000$ and initial condition of 100 infected and the remaining susceptible. The SIR model has an initial fraction $\frac{100}{100,000}$ of the population infected and the rest susceptible. The agreement provides some justification for our use of the mean-field approximation. The figure also provides an example of typical SIR epidemic curves.

Equations (4)-(6) are the typical form of the simple compartmental SIR models encountered in the literature, with $S(t), I(t)$ and $R(t)$ representing the fraction of the population in each disease state. Often, for the sake of brevity, the explicit dependence on t is dropped.

To verify our derivations, we compare simulations of the original discrete-time, discrete-space model with predictions of the derived differential equation model. The black dots in Figure 1 are a plot of 200 runs of the discrete-time, discrete-space model with a population size of $100,000$ individuals, 100 of whom are initially infected and the rest susceptible. The three lines in the figure represent the solution to the numerically integrated compartmental SIR model. For both the discrete-time model and the differential equations, we set $\beta, \gamma = 0.4, 0.2$. The agreement between the discrete model and the differential equation model that we see in Figure 1 provides some justification for the mean-field approximation we used in our derivations.

Figure 1 also provides an example of the typical epidemic curves both seen in real-world epidemics and produced by SIR models. Initially, the epidemic and the number of infected individuals grows exponentially. However, there is a turning point when more infected individuals leave the infected compartment than enter it. The epidemic ends when the number of infected individuals drops to 0, which often happens before all susceptible individuals in the populations are infected.

We can now ask simple, yet illuminating questions of the compartmental SIR model. Perhaps the first and most important question to ask is: Which diseases have the ability to spread in the population and thus become epidemics? SIR models parameterize diseases using two parameters: the infectivity parameter, β, and the infectious period parameter, γ. We can restate the basic question as: For what values of β and γ will we see an epidemic?

Intuitively, an epidemic grows when an infected individual, throughout their entire infectious

period, creates more than one newly infected individual. For example, suppose each infected individual creates two new infections throughout their infectious period. If we start with only a single infected individual, when that first person recovers from the disease, there will be two new infected individuals. When those two individuals recover, there will be four new infected individuals, and so forth. This leads to the exponential growth of the epidemic.

In the SIR model, individuals leave the infected compartment at a rate γ, giving an infectious period of $\frac{1}{\gamma}$ for each individual. At the beginning of the epidemic, when $S(t)$ is close to 1 and $I(t)$ is just above zero, each infected individual creates new infections at a rate of β. So, the total number of new infections created by each infected individual throughout their entire infectious period is $\frac{\beta}{\gamma}$. Thus, from our intuitive derivation, for the SIR model, if $\frac{\beta}{\gamma} > 1$, a disease will become an epidemic.

The same result can be derived mathematically. We simply want $\frac{dI(t)}{dt}$ to be greater than 0 at the onset of the outbreak. If $S(t)$ is close to 1 and $I(t)$ is just above zero, we can state the condition as $\beta \cdot S(t) \cdot I(t) - \gamma \cdot I(t) \approx \beta \cdot I(t) - \gamma \cdot I(t) > 0$. Which gives $\beta \cdot I(t) > \gamma \cdot I(t)$, or $\frac{\beta}{\gamma} > 1$.

To address the fundamental question of "Which diseases become epidemics?", we define a useful epidemiological quantity. Let R_0, also called the *basic reproduction number*, be the expected number of new infections created by an infected individual under the most favorable conditions for transmission. For the SIR model, we have $R_0 = \frac{\beta}{\gamma}$. In general, for any disease in any host population, the disease can become an epidemic only if $R_0 > 1$. The mathematical condition $R_0 > 1$ can be intuitively interpreted as saying that there exist some conditions under which the disease can grow. For the SIR model, those "most favorable conditions" are when $S(t)$ is close to 1 and $I(t)$ is just above zero. Since R_0 describes the number of new infecteds created by each infected individual, then, during the earliest stages of an epidemic, the number of infected individuals in the ith generation of transmission is roughly R_0^i.

The behavior of the basic SIR model varies as we alter R_0, β and γ. Figure 2(a) depicts the infected compartment curves as we fix γ to 0.2 and alter β, and thus R_0. As can be intuitively expected, values of R_0 that are close to one produce very slow-growing epidemics, while values of R_0 much greater than one produce fast, explosive epidemics.

If we were to keep R_0 fixed, but alter γ, we could stretch any of the curves in Figure 2(a) along the horizontal axis. Intuitively γ provides a time-scale of the epidemic. However, the final epidemic size is fixed by the value of R_0. This is demonstrated by Figure 2(b), which plots resistant curves when we keep R_0 fixed, but vary the values of β and γ. For a mathematical derivation of final epidemic size in an SIR model as a function of R_0, see Keeling and Rohani [29].

Finally, though we have discussed a simple SIR model, the approach can be extended to model more complex disease progression, as well as more complex population structures. For example, Equations (7)-(9) represent a model where we have introduced a natural birth/death process that removes individuals from all compartments and introduces individuals into the susceptible compartment. In this model, μ is the birth/death rate. On the other hand, Equations (10)-(13) have introduced a latent period of the disease, between the susceptible and infected compartments. In this model, E stands for the *exposed* compartment of infected individuals in the latent stage of infection; and σ is the rate at which these individuals progress to the active stage of infection. Introducing more complex population structure can also be done to some extent by adding a set of SIR variables for each group of individuals in the population. See Keeling and Rohani [29] for more on compartmental models.
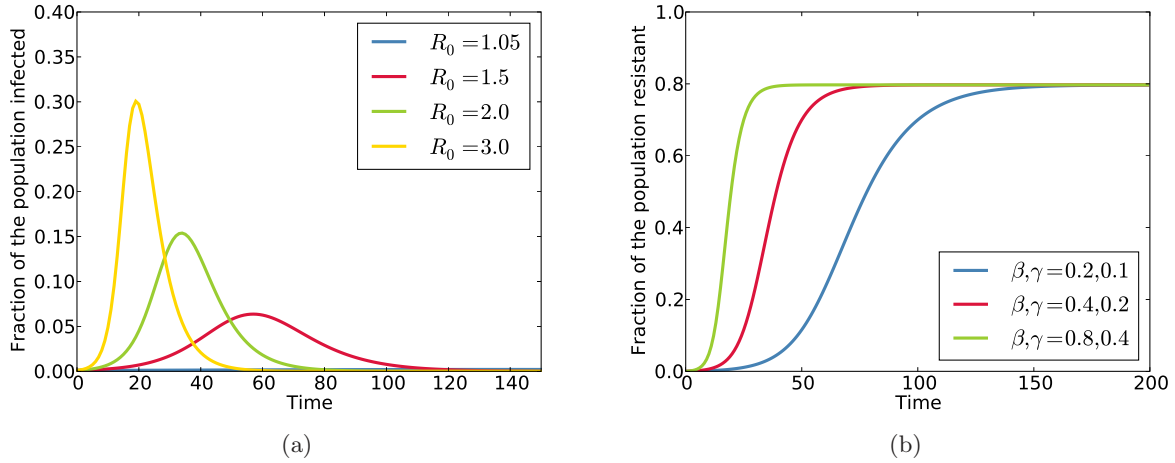
Figure 2: Dynamics of SIR compartmental models. (a) The rate and magnitude of the epidemic varies with R_0. The y-axis shows the fraction of the population currently infected. Here γ is fixed to 0.2 and R_0 varies. Diseases with R_0 close to one produce slow-growing epidemics. Higher values of R_0 yield quickly growing, explosive epidemics. (b) The progression of the epidemic varies with γ and β. The y-axis shows the changing fraction of the population in the resistant compartment. Here R_0 is held constant while β and γ vary. The infectious period parameter γ provides a time-scale for the epidemic, hastening the epidemic as γ increases, which decreases the infectious period. However, the final epidemic size is fixed by the constant value of R_0.

$$\frac{dS}{dt} = -\beta \cdot S \cdot I + \mu \cdot (I + R) \qquad (7)$$

$$\frac{dI}{dt} = \beta \cdot S \cdot I - \gamma \cdot I - \mu \cdot I \qquad (8)$$

$$\frac{dR}{dt} = \gamma \cdot I - \mu \cdot R \qquad (9)$$

$$\frac{dS}{dt} = -\beta \cdot S \cdot I \qquad (10)$$

$$\frac{dE}{dt} = \beta \cdot S \cdot I - \sigma \cdot E \qquad (11)$$

$$\frac{dI}{dt} = \sigma \cdot E - \gamma \cdot I \qquad (12)$$

$$\frac{dR}{dt} = \gamma \cdot I. \qquad (13)$$

3 Estimating Epidemiological Rates and Constants

Epidemiological models are only as good as their parameter values. That is, accurate forecasting and understanding of disease dynamics requires finding and using realistic epidemiological rates and constants in the equations. As we have already seen, the value of R_0 crucially affects epidemic dynamics. In addition, inaccurate estimates of the infectious period can lead to miss-calculating the timing of epidemic peaks, when resources are most needed. Thus, much research effort is devoted to

accurate parameter estimation of emerging infectious diseases. In this section, we briefly describe some of the methods and pitfalls in estimating disease parameters using real-life data.

Typical data available at the beginning of an epidemic includes a time-series of new-case occurrences, from surveillance, physician reports, or hospitalizations [28]. Occasionally, data describing individual-to-individual chains of transmissions is also available. For example, a family member could contract the disease abroad and return to infect other family members some time later. In a recent example, an index case of H1N1/09 created 13 new cases during a bus trip on the way to a soccer match in Scotland [53].

Parameterization often starts with estimation of the time-scale and rate of growth of the epidemic. There are two dominant approaches to estimating the time-scale of transmission. Using the first approach, one can estimate the *generation time*, defined as the expected length of time between infection of an index case and infection of his/her secondary cases [51]. Inter-infection times between the index case and all those whom they infect are simply averaged. However, determining the timing of infections can be complicated for diseases with asymptomatic periods of unknown or variable duration. Using the second approach, one can estimate a quantity closely related to the generation time, the *serial interval*, defined as the time between the clinical onset of symptoms in the index case and the clinical onset of symptoms in the average secondary case. Both approaches require data on individual-to-individual chains of transmission.

The rate of growth of an epidemic is typically estimated from time-series data of new cases. The basic reproduction number (R_0) gives us the ratio between the numbers of infected individuals in consecutive generations of the epidemic. If the generation time has already been estimated, then the time-series of new case occurrences can be grouped into generations. One can then estimate R_0 by fitting an exponential growth to the resulting grouping, using an appropriate statistical model [11, 13, 35]. Recently, higher fidelity methods, accounting for non-homogeneously mixed populations and delays in reporting, have been developed for estimating R_0 [47, 58, 62]. In Appendix B, using an SIR model, we provide an example of the complex and model specific methods required for model parameterization.

4   Limitations of Compartmental Models

While compartmental SIR models have proven to be quite useful in modeling epidemics, they do not properly model some important aspects of disease spread. For example, consider the 2002-2003 outbreak of Severe Acute Respiratory Syndrome (SARS). Estimates of R_0 based on the initial outbreak of SARS ranged between 2.2 and 3.6 [35, 50]. The case fatality ratio was estimated to be between 11% and 13% [18, 63]. For comparison, the U.S. Department of Health and Human Services assigns the greatest pandemic severity ranking to pandemics with a case fatality ratio of 2%; pandemics with this ranking would require the strictest national response strategies [54, 52]. Based on the estimates of R_0, SARS should have caused a great world pandemic with cases numbering easily in the millions. However, for the entire SARS outbreak (from November 1st 2002 to July 31st 2003), only 8,096 cases were reported with 774 deaths [64].

Certainly, one explanation for the limited spread of SARS is the quick response by world public health agencies, who imposed strict quarantines on infected individuals. Another likely explanation for the discrepancy is that the estimates for R_0 were based on data involving large numbers of transmissions in hospitals, where people have unusually high rates of contact. SIR models assume a fully mixed, homogeneous population – the *mass-action* assumption – in which each individual has the same amount of contacts as every other individual. Thus, simple SIR models do not accurately model the increased rate of contact at hospitals and the decreased rate of contact of quarantined
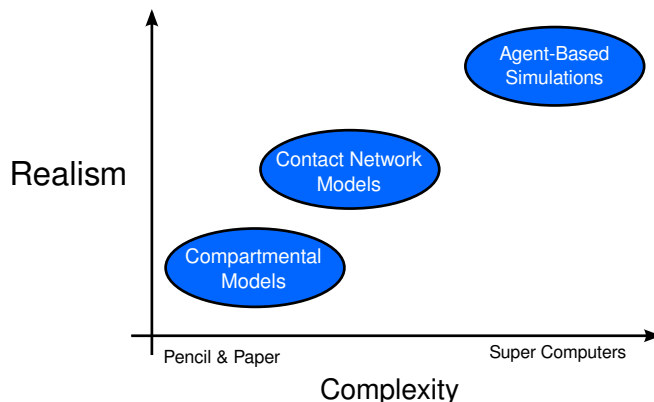
Figure 3: The complexity of epidemiological models. It is often useful to think of models in two dimensions: the extent to which they capture real-world complexities that impact disease transmission (y-axis) and their computational tractability (x-axis). Compartmental models are easy to analyze but miss important, realistic details, such as heterogeneous patterns and types of contacts. Agent-based simulations are able to model reality with a great amount of detail, but are difficult to parameterize, analyze, and require large amounts computation. Contact network models capture disease transmission with a higher fidelity than compartmental models, yet remain analytically tractable.

individuals. If the population at large had as many contacts as the population within a hospital, perhaps the estimates of R_0 would have been more accurate, and SARS would have infected many more people.

Incorporating realistic contact patterns of the population is just one possible way to increase the fidelity of epidemic models. Diseases often spread differently in different age groups, have varying incubation periods in different age groups, spread differently depending on the type of contact – e.g., contacts at home tend to be more intimate than contacts at work. Also, disease spread is affected by geographic location and seasonality. Researchers have built very high-fidelity models using agent-based simulations, where each individual is tracked as they move from home, to work and back [19]. Naturally, such models involve complex parameterization and often require extensive computation (see Figure 3). In the next section, we introduce contact network models, a type of epidemiological model that lies between between compartmental models and agent-based simulations, providing higher fidelity, yet tractable formulations.

5   Contact Network Modeling

Contact network epidemiology is an analytical framework that intuitively captures the diverse host interactions that underlie parasite transmission [39, 45]. The first step in this modeling approach is to build a realistic network model of contact patterns at an appropriate temporal and spatial scale. The second step is to predict the spread of disease through the resulting network, based on intrinsic features of the parasite and the network structure. To mathematically analyze disease spread through networks, we apply generating function methods adapted from an area of statistical physics called percolation theory [27].

5.1 Building a contact network model

A contact network model uses a graph to capture patterns of interactions that can lead to parasite transmission. Each host (or group of hosts) translates into a node and contacts among hosts (or groups) translate into edges connecting appropriate nodes. The number of edges emanating from a node is called the *degree* of the node and indicates the number of contacts along which parasite transmission is possible. The distribution of the number of such contacts within a population, called the *degree distribution*, fundamentally influences the spread of pathogens through the population.

The study of contact networks, and more generally of social networks, is of growing importance in a diverse group of disciplines [3, 61, 31]. Researchers seek universal properties, and have focused on *small-world* networks – with high levels of both local clustering and global connectivity [60] – and *scale free* networks – with power-law degree distributions leading to a small fraction of very highly connected hubs [10].

Several epidemiologically relevant networks including sexual contact networks have been characterized as scale free [34, 46]; however, researchers have found that realistic contact networks do not always exhibit these well-studied structural properties [7]. For example, contact networks underlying the spread of respiratory and air-borne diseases tend to have degree distributions that appear more exponential in shape than scale free, homogeneous (all nodes have the same degree), or random (Poisson degree distribution) (see Figure 4). In the remainder of this section, we describe a various approaches to constructing realistic contact networks and provide three specific examples from the literature.

To construct a contact network for any particular disease or class of diseases, we first define an epidemiological contact. For respiratory diseases, this may mean close physical proximity for a specified duration; or, for sexually transmitted diseases, this may mean having sexual relations or sharing needles. We then seek data on the distribution of such contacts across the focal population. This data may come from sociological surveys [43, 49] or wildlife studies [15]; or the distributions can be inferred from general information about activity patterns, using statistical tools or computer simulations that generate explicit networks from such patterns [19, 42].

On a small scale, we have modeled the activity patterns that underlie the transmission of respiratory-borne diseases in a psychiatric institution in Evansville, Indiana based on detailed information about the distribution of caregivers and patients among wards [40]. We represent these contacts in a bipartite network. One set of nodes represent health care workers and the other represent entire wards filled with patients (see Figure 5(b)). For this particular institution, there were few if any direct contacts between patients in different wards or between caregivers outside of wards, and thus we ignore ward-ward and caregiver-caregiver contacts in the model.

On an intermediate scale, we have developed software to generate contact network models for an urban setting [42, 48, 9]. Based on detailed demographic, employment, school, and hospital data for the city of Vancouver, British Columbia, we model interactions within homes, schools, neighborhoods, work, hospitals, shops, restaurants, etc. We start with up to one million households, drawn at random from the Vancouver household size distribution, which yields up to 2.6 million people. Household members are assigned ages according to the Vancouver age distribution. Each individual, based on age, is then assigned to daycare centers according to early childhood care statistics, to schools according to school and class size distributions, to occupations according to employment data, to hospitals as patients and caregivers according to hospital employment and bed occupancy data, and to other public places. Within each location we create random connections between individuals with probabilities ranging from zero to one, depending on the type of location.
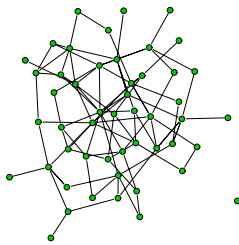
The resulting network is undirected, meaning that transmission may occur in either direction along an edge (see Figure 5(a)). For example, two individuals in the same household will have equal
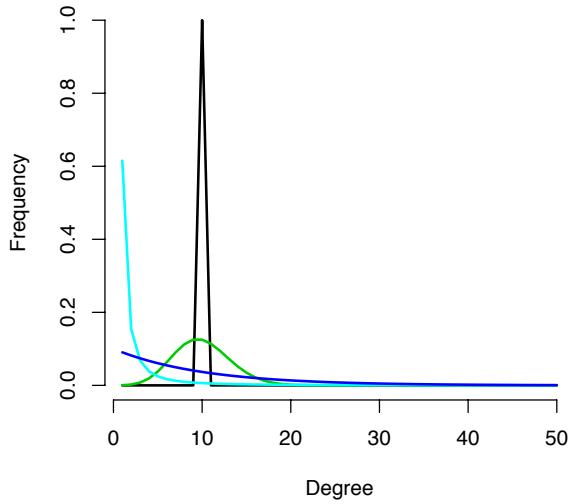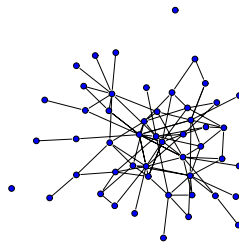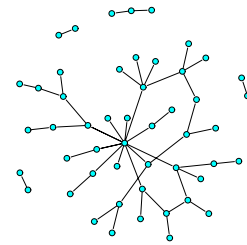
(a) Degree distributions
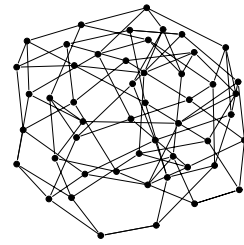
(b) Homogeneous network

(c) Poisson network

(d) Exponential network

(e) Scale-free network

Figure 4: Comparing networks with different degree distributions. (a) Several degree distributions with a mean degree of 10: homogeneous, where all nodes have exactly 10 contacts (black), Poisson distribution (green), exponential distribution (blue), power law distribution (cyan). (b)-(e) Examples of homogeneous, Poisson, exponential, and scale free networks, respectively, with degree distributions shown in (a). Contact networks underlying spread of respiratory diseases like flu tend to have exponential-like distributions (blue).

opportunities to infect each other. There are cases, however, where a person may infect another person but the converse is not true. Suppose individual A is normally healthy and thus has no reason to go to the hospital until he or she becomes infected with a severe infectious disease. At that point, individual A may come into contact and potentially spread disease to health care workers (HCW). In contrast, if an HCW acquired the disease while individual A remained healthy, then there would be no opportunity for transmission in the opposite direction. To model the unidirectional flow of disease into hospitals we include directed edges from individuals in the population at large to HCWs, yielding a *semi-directed* network with both directed and undirected edges (see Figure 5(c)). For diseases like SARS and severe pandemic flu, most individuals would likely seek medical treatment upon developing symptoms, and thus the contact network model for these diseases would include directed edges from most of the population to HCWs. In contrast, for a disease like seasonal flu, only the high-risk populations – such as the very young, elderly, or immunocompromised – tend to seek hospital care, and thus the seasonal flu contact network model would include directed edges for high-risk groups only.

On a large scale, we have modeled the connectivity among the largest cities in North America. In this case, the nodes are cities and the edges reflect travel patterns between cities via air and ground transportation, as reported by the US Census Bureau, the US Bureau of Transportation Statistics, Instituto Nacional de Estadística y Geografía, and Statistics Canada. In this case, the edges of our network are weighted by travel flux and diseases spread within cities via simple compartmental models [17] (see Figures 5(d) and 9(a)).

5.2 Using Bond Percolation to Model Susceptible-Infected-Recovered (SIR) Disease Dynamics

Imagine that a parasite initially appears at a random node in a contact network. The disease propagates through the network similarly as in an SIR compartmental model, except that the spread is guided by the structure of the contact network instead of the uniformly random contact patterns of a compartmental model. The initial node remains infectious for some period of time, during which it has the potential to transmit disease to each of its contacts. The secondary cases likewise can transmit disease to their contacts during their infectious periods, and so on. This process resembles simple bond percolation from statistical physics which models, for example, the flow of a liquid through a porous material [27]. Just as the liquid traverses gaps in the porous material with a characteristic viscosity, a disease spreads from person-to-person with a characteristic level of infectiousness. In general, percolation theory describes connectivity in random graphs, and thus can be applied to predict the size of the infected cluster, that is, the number of nodes reached via parasite transmission along the edges in the network. This approach was initially suggested by Grassberger [26] and Newman [45]. Recently, we have extended it into a flexible framework for infectious disease modeling [40, 42, 41]. These methods allow us to make predictions for infinite networks with a specified degree distribution. To use it, we must assume that (1) the contact network is infinite (or quite large) and (2) the epidemiologically-relevant structure of the network is adequately summarized in its degree distribution. The second assumption means that we ignore additional structure, like local clustering, beyond what is expected in an infinite network with the given degree distribution. To test that these assumptions are reasonable, we often compare our mathematical predictions (based only on the degree distribution of the network) to simulations of disease spread through the full finite-sized contact network.

The fate of an outbreak depends on both the level of contagion and the structure of the underlying contact network. To model contagion, every edge in a network is given a probability of pathogen transmission along it (T_{ij}), that is, the probability that host i, if infected, will transmit

disease to individual j during his or her infectious period. If it is reasonable to assume that the T_{ij} are independently and identically distributed (iid) random variables, then we can make calculations based solely on the average of these probabilities (T) [45]. This value summarizes core aspects of disease transmission including the rate at which contacts take place between hosts, the likelihood that an encounter will lead to transmission, the duration of the infectious period, and individual susceptibility. When the T_{ij} are not iid, percolation calculations are possible but more difficult.

In these calculations, we use the degree distribution of a network to indicate its structure. Probability generating functions (pgfs) are functions that completely describe discrete probability distributions. For infectious disease modeling, the pgf of a contact network's degree distribution summarizes useful information about the structure of the contact network. For example, the pgf for the degree distribution of an undirected random network is $G_0(x) = \sum_{k=1}^{\infty} p_k x^k$, where p_k is the relative frequency of nodes of degree k in the network. Using straightforward probabilistic arguments, we sequentially derive pgfs for the distributions of (1) the number of edges emanating from a node reached along a randomly chosen edge: $G_1(x) = G_0'(x)/\langle k \rangle$, where $\langle k \rangle = G_0'(1)$ is the average degree in the network, (2) the number of edges along which disease transmits from an infected node during an outbreak: $G_0(x;T) = G_0(1+(x-1)T)$, where T is the average probability of transmission, and (3) the size of outbreaks stemming from a single introduction of disease: $H_0(x;T) = xG_0(H_1(x;T);T)$, where H_1 is defined by the self-referential pgf $H_1(x;T) = xG_1(H_1(x;T);T)$ [45].

For a random network with a given degree distribution, there typically exists a threshold transmission rate below which small, finite-sized outbreaks occur and above which large-scale epidemics, comparable to the size of the network, are possible. This *epidemic threshold* is analogous to the well-studied *percolation threshold*[1], and it depends on the network structure. In an undirected random network with a given degree distribution, for example, the epidemic threshold is $T_c = 1/G_1'(1)$. Intuitively, $G_1'(1)$ can be interpreted as the pgf for the number of susceptible contacts for each infected individual. So, as in our discussion of the basic reproductive number, an epidemic occurs if, in expectation, each infected individual creates more than one newly infected individual, i.e. if $T \cdot G_1'(1) > 1$. Highly connected networks, with ample opportunities for transmission, have low epidemic thresholds. In such networks, even mildly transmissible parasites will be able to cause epidemics. Less connected networks will have higher epidemic thresholds. The pgf approach also allows us to compute the expected size of an outbreak for a pathogen below the epidemic threshold, $1 + \frac{TG_0(1)}{1-TG_1'(1)}$, and both the probability and expected size of a large epidemic for a parasite above the epidemic threshold, which are equal to each other in undirected networks, $1 - G_0(u;T)$ where the self-referential $u = G_1(u;T)$ can be solved numerically [45].

Newman introduced the pgf approach to analyzing epidemics [45]. We have since derived similar quantities for bipartite and semidirected networks [40, 41], and extended it to calculate a number of important epidemiological quantities on networks, including: (1) epidemic threshold, (2) expected size of a small outbreak, (3) probability of a large-scale epidemic, (4) expected size of a large-scale epidemic (should one occur), (5) quantities 1-4 conditioned upon the identity of the node (or nodes) where the parasite first appeared, (6) quantities 1-4 conditioned upon the size of an initial outbreak, (7) the probability that a specific node will become infected during an epidemic, and (8) the degree distribution of the residual network (the remaining network of uninfected nodes following an epidemic) [40, 42, 41, 21, 8].

---

[1]In a network in which every pair of nodes are connected with probability p, the percolation threshold is the value of p above which connected clusters are expected to span the entire (infinite) network.

5.3 Dynamic Network Models

The bond percolation approach captures an important aspect of population heterogeneity, but has two important shortcomings. First, it predicts the final state of an outbreak, but not the temporal progression of disease. Second, it assumes that the contact network is static, that is, that the numbers and identities of a node's contacts are fixed throughout the outbreak. Although this assumption may be reasonable for rapidly spreading diseases, there are many situations in which the underlying network will change considerably during an outbreak. For example, concurrent and serial contacts are known to strongly influence the spread of sexually transmitted infections like HIV [59, 22, 1].

Volz recently developed a low-dimensional system of non-linear ordinary differential equations to model the dynamical progression of a disease spreading through static random networks with arbitrary degree distributions [55]; and we have extended this framework from static networks to dynamic networks [56]. This model improves on the bond percolation approach in that it both predicts the temporal progression of disease and allows for changing structure in the underlying network. Specifically, the model considers a simple class of dynamic networks in which pairs of edges are randomly chosen and swapped. For example, if edges AB and CD are chosen, then they are deleted and replaced by edges AD and CB. In this neighbor exchange model, each node maintains a constant number of contacts, but the identities of those contacts may change randomly. The model is given by the following equations:

$$\frac{d\theta}{dt} = -\theta r \frac{M_{SI}}{M_S}$$

$$\frac{dM_{SI}}{dt} = rM_{SI}\left(-1 - \frac{\delta}{M}\frac{M_{SI}}{M_S} + \frac{\delta}{M}\frac{M_{SS}}{M_S}\right) - M_{SI}\mu - \rho(M_{SI} - M_I M_S)$$

$$\frac{dM_{SS}}{dt} = -2rM_{SI}\left(\frac{\delta}{M}\frac{M_{SS}}{M_S}\right) - \rho(M_{SS} - M_S M_S)$$

$$\frac{dM_I}{dt} = rM_{SI}\left(\frac{\delta+1}{M}\right) - \mu M_I$$

The model consists of four core dynamic variables: θ is the fraction of degree one nodes that are still susceptible; M_{SI} is the fraction of edges in the network connecting a susceptible node and an infected node; M_{SS} is the fraction of edges in the network connecting two susceptible nodes; and M_I is the fraction of edges in the network adjacent on an infected node, regardless of the state of the node at the other end of the edge. There are also four fixed parameters: r is the transmission rate; μ is the recovery rate; $g(x)$ is the pgf for the network degree distribution; and ρ is the neighbor exchange rate. To simplify the equations, we also use three helper values: $M = g'(1)$ is the total number of edges in the network; $M_S = \theta g'(\theta)$ is the fraction of edges adjacent on a susceptible node; and $\delta = \frac{\theta g''(\theta)}{g'(\theta)}$ is the average excess degree for a susceptible node selected by following a random chosen $I \leftrightarrow S$ edge. *Excess degree* is defined as the degree of the node minus one. The last two of these helper values vary as the epidemic progresses through the network. Finally, the equations highlight the commonly appearing term rM_{SI}, which is the rate of transmission events in the network per unit time.

This model tracks the state of each edge and each stub (one end of an edge) as disease spreads through the network. Figure 6 illustrates the impact of new infections, recoveries, and neighbor exchanges on the composition of edges in the network. To provide some intuition behind these equations, we deconstruct each one here.

13

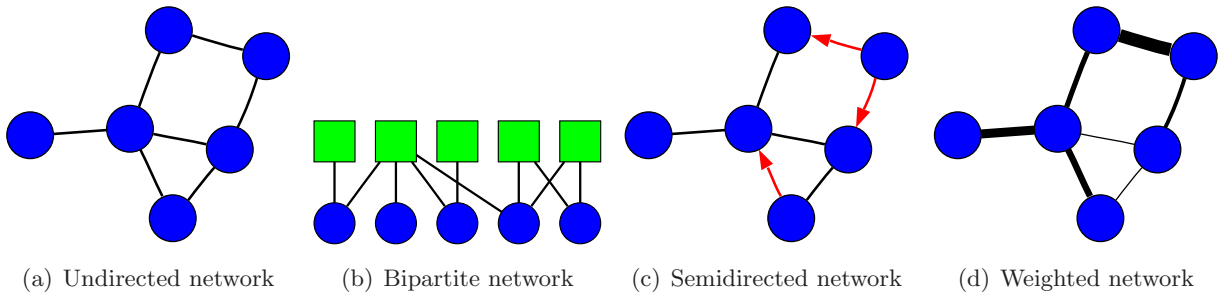(a) Undirected network   (b) Bipartite network   (c) Semidirected network   (d) Weighted network

Figure 5: Common classes of networks used to model disease spread. (a) Simple undirected network. This type of network has been used for modeling person-to-person contacts. (b) Bipartite network. This type of network has been used for modeling contacts between caregivers and wards. (c) Semidirected network. This type of network has been used for modeling the one-way contacts from the general population to heath care workers. (d) Weighted network. This type of network has been used for modeling travel patterns between cities.

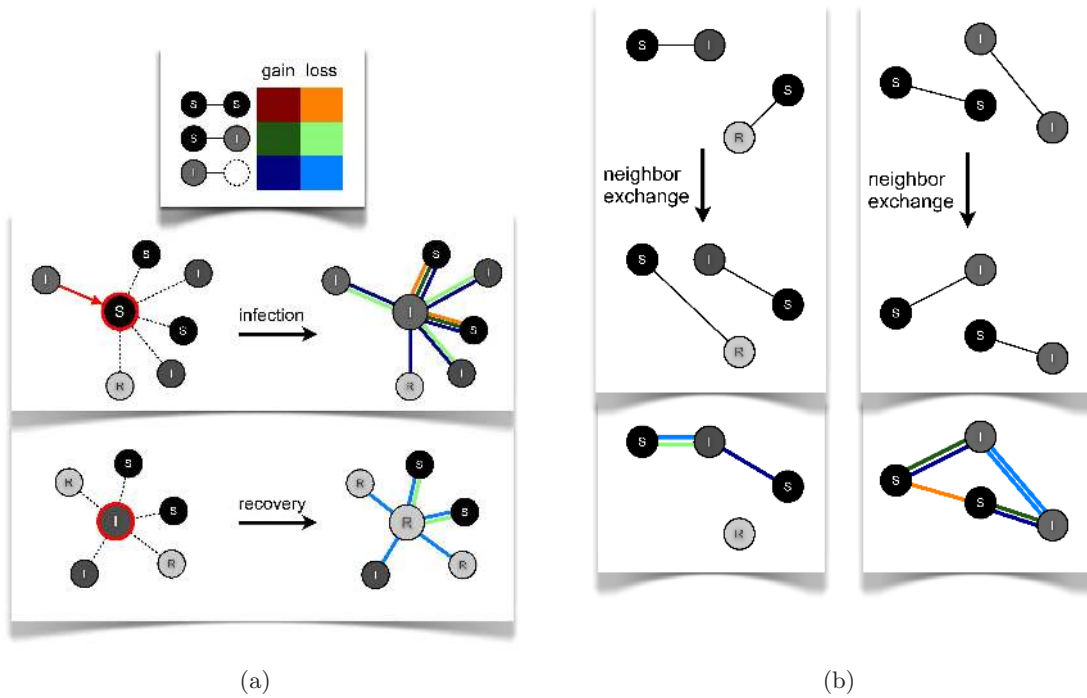(a)   (b)

Figure 6: The impact of infections, recoveries and neighbor exchanges on the composition of edges in a network. Edge colors indicate which types of edges are created and destroyed following each event. (a) New infections and recoveries lead to gains and losses various edge types. (b) Two examples of neighbor exchange events and their impacts on edge composition (below). There are many other ways in which neighbor exchange events can impact edge composition.

14

The first equation describes the decline in the number of degree one nodes that are susceptible. If a degree one individual is susceptible, then $\frac{M_{SI}}{M_S}$ is the probability that his/her single edge is connected to an infected node and r is the probability of transmission along that edge. Thus $r\frac{M_{SI}}{M_S}$ is the rate at which such nodes become infected.

The second equation describes the change in the fraction of edges connecting susceptible nodes to infected nodes. This is illustrated by the gains (dark green) and losses (light green) in Figure 6. Consider one such edge, connecting a susceptible node attached to an infected node. The first term in the equation performs the accounting required due to transmission events. When a transmission event occurs, turning a susceptible node to infected, we must (1) remove the single $S \leftrightarrow I$ edge carrying the infection (2) remove any other $S \leftrightarrow I$ edges adjacent on the newly infected node (3) add any $S \leftrightarrow S$ edges adjacent on the newly infected node. The second term accounts for recovery events, which convert edges from $S \leftrightarrow I$ to $S \leftrightarrow R$. The final term corresponds to the impact of neighbor exchanges on the fraction of $S \leftrightarrow I$ edges. By randomly mixing the network edges, neighbor exchanges slowly bring the fraction of $S \leftrightarrow I$ edges to the expected fraction of such edges, $M_S M_I$. The value $M_S M_I$ corresponds to the expected number of $S \leftrightarrow I$ edges in a network that has the same fraction of stubs connected to susceptibles and infecteds as the original network, but has edges re-distributed randomly between nodes.

The third equation describes the change in the fraction of edges connecting susceptible nodes to other susceptible nodes. This is illustrated by the gains (red) and losses (orange) in Figure 6. Consider one such edge connecting a susceptible with another a susceptible node. The first term in the equation corresponds to loss of $S \leftrightarrow S$ edges following an infection transmission event. The $S \leftrightarrow S$ edges that are turned into $S \leftrightarrow I$ edges, added in the previous equation, must be subtracted here. The rate of change on $S \leftrightarrow S$ edges is doubled, since both their endpoints become infected at the same rate. As before, the final term corresponds to the impact of neighbor exchanges on the fraction of $S \leftrightarrow S$ edges. Neighbor exchange acts like a spring with tension ρ slowly bringing the fraction of $S \leftrightarrow S$ edges to the expected value $M_S M_S$ for a comparable random network.

The final equation describes the change in the fraction of stubs that are adjacent on an infected node. This is illustrated by the gains (dark blue) and losses (light blue) in Figure 6. The first term performs accounting due to newly created infections. When a susceptible node becomes infected all of the edges emanating from the node add to the class of M_I stubs, including the single edge involved in the transmission. The second term in the equation governs the loss of infected stubs through the recovery of infected nodes.

This model fares well in comparison to stochastic simulations of an analogous epidemic process in networks, that is, it predicts an epidemic trajectory (cumulative incidence curve) that falls right in the middle of the curves produced by stochastic simulations (see Figure 7(a)) and makes good predictions for the final state of the epidemic (cumulative number and distribution of cases) [56, 57].

5.4 Connecting Compartmental and Network Models

The Volz-Meyers dynamic network model is not only very tractable (with just one more dynamic variable than the standard SIR compartmental model), but also offers a mathematical and conceptual bridge between two disparate classes of models. That is, by changing the value of the mixing parameter ρ, we interpolate smoothly between models without neighbor exchange ($\rho = 0$) and compartmental models ($\rho = \infty$). In the limit of large mixing ($\rho \to \infty$), every transmission event from infected nodes is essentially directed at a randomly chosen node. That is, the probability of being connected along any given edge to a susceptible, infectious, or recovered node is directly proportional to the number of edges connected to nodes in each of these states, respectively. The resulting model is thus a mass-action model with three dynamic variables that allows arbitrary

heterogeneity in contact rates, as quantified by the pgf $g(x)$. If we assume that contact rates are homogeneous ($g(x) = x$), then the model exactly reduces to the standard SIR compartmental model described earlier (see Figure 7(b)).

5.5 Advantages and Limitations of Network Models

The basic models described here make many simplifying assumptions about host population structure and epidemiological parameters. For example, the population structure is assumed to resemble a graph with a specific degree distribution and transmission and recovery rates are assumed to be homogeneous across both hosts and time. In the last few years, however, these models have been extended to incorporate additional complexity, including dynamic contacts, assortative connectivity, and heterogeneity in transmission rates [45, 56, 44, 8]. While each model has its own limitations, the framework as a whole has been shown to be versatile and is evolving to address more complex ecological data and questions.

Since the models are simpler than the real populations they represent, modelers typically check mathematical predictions through comparisons to actual ecological data and the results of more complex agent-based simulations [48, 21, 9, 56]. In most cases, the analytical predictions are consistent with the data. When significant discrepancies arise, modelers work to identify and incorporate key dynamics missing from the model.

Although there are now many sophisticated mathematical approaches to modeling host-pathogen dynamics, the network methods described here have several advantages. Like agent-based models and other individual-based models, they have the advantage over compartmental models of simply and intuitively capturing heterogeneity in host contact patterns. Contact heterogeneity profoundly influences host-pathogen dynamics both quantitatively and qualitatively and ignoring network structure can lead to erroneous predictions [42, 56, 7]; and even small quantitative differences can be critical to effective public health and environmental management. Contact network models have already provided important insights into disease dynamics. For example, they have shown prior outbreaks of an immunizing disease like influenza can dramatically influence the dynamics of future outbreaks, although the impact depends on the structure of the host network [21, 8]. This occurs because disease preferentially infects the most highly connected demographics. Contact network models have also shed light on the heterogeneous spread of SARS [42] and the role of hospitals in community outbreaks [41] and have been use to design effective control strategies for respiratory diseases in health care and urban settings [48, 40] and optimal vaccination strategies for influenza [9, 8]. Another important advantage of contact network methods is that they are mathematically simpler than agent-based models for capturing heterogeneous contact patterns, which allows for rapid and accurate calculations and the derivation of analytical results.

6 Disease Control

Infectious disease models help us not only understand the dynamics of spreading pathogens but also design effective strategies for controlling outbreaks. In this section, we describe some of the primary modes of infectious disease intervention and illustrate how mathematical and computational methods can be used to optimize such interventions.

6.1 The fundamentals of disease control

Consider the SIR model described by Equations (7)-(9), which incorporates the natural birth and death of individuals into the simple SIR model. To control an epidemic at any point, we would

like to decrease $\frac{d\mathrm{I}}{dt}$, the number of new infected individuals created per unit time. If we are able to make $\frac{d\mathrm{I}}{dt}$ negative, then the number of infecteds will begin to decrease. When the number of infecteds reaches 0, the epidemic will end.

What actions can we perform to make $\frac{d\mathrm{I}}{dt}$ negative? Using Equation (8), we have

$$\frac{d\mathrm{I}}{dt} = \beta \cdot \mathrm{S} \cdot \mathrm{I} - \gamma \cdot \mathrm{I} - \mu \cdot \mathrm{I} < 0.$$

With rearranging terms and dividing by I, the above expression becomes

$$\frac{\beta \cdot \mathrm{S}}{\gamma + \mu} < 1. \tag{14}$$

The left hand side of Equation (14) is called the *effective reproduction number*, or effective R, at the current point of the epidemic. The effective R gives us an idea of how quickly the epidemic is currently growing. If we can get the effective R bellow one, the epidemic will begin to die. According to Equation (14), to reduce the effective R, and thus control the epidemic, we have the following options:

1. Reduce β, the infectivity parameter.
2. Reduce S, the fraction of susceptibles in the population.
3. Increase γ, the infectious period parameter.
4. Increase μ, the natural death rate of individuals.

Let us take each of these options in turn and give them realistic interpretations. The infectivity parameter β can be thought of as the product of the likelihood the disease is transferred during contact, and the likelihood that contact occurs between an infected and susceptible individual. Thus, β can be reduced by actions like (1) quarantining infected individuals, reducing use of public transport, closing schools, or encouraging the workforce to work from home, all of which reduce the likelihood of contacts between infecteds and susceptibles, (2) increasing hand washing and other hygienic precautions that potentially reduce the likelihood of transmission during contact, and (3) rapidly treating infected individuals with antimicrobials may reduce symptoms that would otherwise enhance transmission during contacts.

Reducing S, the number of susceptibles in the population, through vaccination is a critical and often long-lasting disease control strategy. When susceptible individuals are effectively immunized, they move to the resistant compartment without experiencing the disease. Some vaccines provide immunity that lasts for decades or even a life-time, and thereby severely limit the potential for future transmission. Equation (14), specifying the effective R, can be used to derive the fraction of the population that must be vaccinated to prevent future growth of the epidemic. Specifically, if the fraction of susceptibles in the population is reduced to less than $\frac{\gamma + \mu}{\beta}$, then the disease is unable to spread. This example demonstrates how partial vaccination of a population, reducing S to a small but non-zero value, is sufficient to protect the population as a whole. This phenomenon is called *herd immunity*.

The infectious period, given that the infected person does not experience a natural death, in the model under consideration is $\frac{1}{\gamma}$. Increasing γ, the infectious period parameter, is the same as decreasing the infectious period. For some diseases, this can be accomplished through treatment with antimicrobials that speed up recovery.

Finally, one can also increase the natural death rate μ. While this is not an ethical option for human outbreaks, it is a strategy often used to control epidemics in livestock. For example, the United Kingdom has used the culling of cows to control foot-and-mouth and mad cow disease [2, 33].

6.2 A network modeling perspective on disease control

In the previous section we discussed strategies for bringing the reproduction number below one. From the perspective of network models, this is equivalent to bringing the average transmissability of a disease T below the epidemic threshold T_c. This can be achieved through interventions that either directly reduce the infectiousness of the pathogen (i.e., lower T), modify contact patterns so that the pathogen cannot easily spread through the population (i.e., increase T_c), or immunize segments of the population (i.e., increase T_c). We call these three forms of intervention *transmission reducing*, *contact reducing*, and *immunizing*, respectively [48].

Transmission reducing interventions introduce physical barriers to interrupt the spread of respiratory droplets or other infectious particles (e.g., face masks, hand hygiene, disinfection of animate objects, or, in the case of sexually transmitted infections, condoms). These interventions can be modeled by reducing the T_{ij}, the probability of transmission from node i to node j, along the corresponding subset of edges.

Contact reducing interventions include isolation of infected persons, quarantine of persons during their incubation period, patient cohorting in hospitals, and closing schools or other public spaces. They can be modeled by removing edges corresponding to contacts avoided. For example, one can model school closures in an urban contact network by removing all edges that represent contacts among students and staff that would occur during school.

Immunizing interventions include prophylactic treatment with antimicrobials and diverse vaccination strategies including ring vaccination (targeting close contacts of current cases), vaccination of priority groups based on risk factors such as age, health, and place of employment, and universal vaccination. Vaccination prior to an outbreak can be modeled by removing nodes from the network corresponding to the effectively immunized individuals.

One can manipulate contact network models to represent a variety of control measures and then use the mathematical methods described above to predict the impact of such measures on disease dynamics. For example, this approach has been used in collaboration with public health officials in the US and Canada to improve public health strategies for controlling walking pneumonia [40] and SARS [48] and for distributing limited supplies of seasonal and pandemic influenza vaccines [9, 8] (see Figure 8).

6.3 Optimizing disease control policies

While predictive models of infectious disease dynamics have become quite powerful, the computational complexity of these models often impedes the systematic optimization of disease control strategies, that is, finding the optimal demographic, spatial and temporal distribution of costly public health resources. Thus, one typical approach within the infectious disease community has been to evaluate a relatively small set of candidate strategies [37, 20, 9, 24], rather than consider the full spectrum of policy options. Recently, however, researchers from diverse fields are beginning to effectively couple infectious disease models with a variety of tools from operations research including simulation optimization techniques [6]. For an extensive review of such methods, see [36].

In this section, we present one recent approach to searching large sets of infectious disease intervention policies [17]. As we describe here, the method is parallelizable, scalable, usable in real time, and can be adapted to work with diverse epidemic models.

Let the sequence $A_1, A_2, \ldots, A_D$ describe a disease control strategy. Though the sequence can be simply an arbitrary string of bits, it may be helpful to think of it as a sequence of actions taken over a specific time period. For example, A_1 describes the control action performed in the first month, A_2 describes the control action performed in the second month, and so forth.

Suppose that we have an infectious disease simulator, $\text{sim}(A_1, A_2, \ldots, A_D)$, that returns the outcome of an epidemic under a specified control strategy. Since disease progression is stochastic, we assume that $\text{sim}(A_1, A_2, \ldots, A_D)$ is a random variable; and we can only sample from its distribution. Further assume that the simulator always returns a real number in the interval $[0, 1]$. For example, the simulator could return the fraction of the population that has not been infected by the end of the epidemic. We can then formulate the disease control problem as follows:

$$\max_{A_1, A_2, \ldots, A_D} E[\text{sim}(A_1, A_2, \ldots, A_D)].$$

To computationally address the above optimization problem, we have designed and implemented the Disease Control System (DiCon) [17, 25]. DiCon is a modular and extensible optimization platform specialized to disease control with the following features:

1. *Swappable, extensible optimization algorithms.* DiCon includes classic algorithms like exhaustive search and more recent algorithms like bandit-based search algorithms [16, 32, 38, 14]. The system also has a simple optimization algorithm interface, so that new algorithms could be added.
2. *Automatic parallelization.* DiCon can both run on your laptop and on supercomputers. With automatic job queueing and processor management, DiCon can handle multiple concurrent optimizations with multiple processors per optimization, scaling to hundreds of processor cores. DiCon manages all communication between processes.
3. *Simple interface.* DiCon has a simple interface. You specify two functions: simulate(), the disease simulator that takes as input a sequence of control actions, and next_action(), which specifies the space of control policies.
4. *Language independence.* DiCon allows you to specify your simulator and space of control policies in any programming language you choose. DiCon uses Google Protocol Buffers to communicate with your stand-alone program. Libraries implementing Google Protocol Buffers are available for most languages, including Python, C++, and Java.
5. *Versatile job specification.* DiCon allows you to run many optimizations with a single command. With a versatile specification language, one can easily try optimizations with varying parameterizations, optimization algorithms, or control policy spaces.
6. *Logging and checkpointing.* DiCon has includes logging capabilities, so that you can track optimizations during computation. DiCon also includes checkpointing, allowing you to resume interrupted optimizations.

A preliminary version of DiCon has been used to compute release schedules for the U.S. National Antiviral Stockpile with the purposes of delaying an influenza epidemic [17]. In that application, the sequence of actions $A_1, A_2, \ldots, A_D$ describes a schedule of antiviral release, with A_i describing both the amount of antivirals to be released in the ith month and the prioritization of those antivirals. The simulator $\text{sim}(A_1, A_2, \ldots, A_D)$ is stochastic, combining a contact network between cities and a compartmental model within each city. The objective of the optimization is to minimize the number of infected individuals in the first year of the epidemic, under the constraint of using no more antivirals than those available in the national stockpile. We used a bandit-based search algorithm to find near-optimal distribution schedules [32, 14, 16]. Even under high levels of loss of released antivirals, through misallocation or misuse, the optimization is able to find effective release schedules to delay the epidemic. Some of the key results of the optimization are summarized in Figure 9. For more details, see Dimitrov et al. [17]

References

[1] A. Adimora, V. Schoenbach, and I. Doherty. HIV and african americans in the southern United States: sexual networks and social context. *Sexually Transmitted Diseases*, 33(7):S39–S45, 2006.

[2] J. Allcock. Why animals have to be slaughtered. The Times (London), February 27 2001. Home News section.

[3] L. A. N. Amaral and J. Ottino. Complex networks - augmenting the framework for the study of complex systems. *European Physical Journal B*, 38(2):147–162, 2004.

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

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

[6] F. Azadivar. Simulation optimization methodologies. In *Proceedings of the 31st ACM Conference on Winter Simulation*, pages 93–100, New York, New York, 1999. ACM.

[7] S. Bansal, B. Grenfell, and L. A. Meyers. When individual behavior matters: Homogeneous and network models in epidemiology. *Journal of the Royal Society Interface*, 4(16):879–891, 2007.

[8] S. Bansal, B. Pourbohloul, N. Hupert, B. Grenfell, and L. A. Meyers. The shifting demographic landscape of influenza. *PLoS One*, 5(2):e9360, 2010.

[9] S. Bansal, B. Pourbohloul, and L. A. Meyers. Comparative analysis of influenza vaccination programs. *PLoS Medicine*, 3(10):e387, 2006.

[10] A. L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[11] N. Becker. Estimation for an epidemic model. *Biometrics*, 32(4):769–777, 1976.

[12] D. Bernoulli and S. Blower. An attempt at a new analysis of the mortality caused by smallpox and of the advantages of inoculation to prevent it. *Reviews in Medical Virology*, 14(5):275–288, 2004.

[13] P. Y. Boëlle, P. Bernillon, and J. C. Desenclos. A preliminary estimation of the reproduction ratio for new influenza A(H1N1) from the outbreak in Mexico, March-April 2009. *Eurosurveillance*, 14(19):pii=19205, 2009.

[14] P. Coquelin and R. Munos. Bandit algorithms for tree search. In *Proceedings of the 23rd Annual AUAI Conference on Uncertainty in Artificial Intelligence*, pages 67–74, Vancouver, British Columbia, 2007. AUAI Press.

[15] M. E. Craft, E. Volz, C. Packer, and L. A. Meyers. Distinguishing epidemic waves from disease spillover in a wildlife population. *Proceedings of the Royal Society of London B*, 276(1663):1777–1785, 2009.

[16] E. Dar, S. Mannor, and Y. Mansour. PAC bounds for multi-armed bandit and markov decision processes. *Proceedings of the 15th Annual ACL Conference on Computational Learning Theory*, pages 255–270, 2002.

[17] N. B. Dimitrov, S. Goll, N. Hupert, Pourbohloul B., and L. A. Meyers. Optimizing tactics for use of the U.S. antiviral strategic national stockpile for pandemic (H1N1) influenza, 2009. *PLoS Currents: Influenza*, page RRN1127, 2009.

[18] C. A. Donnelly, A. C. Ghani, G. M. Leung, A. J. Hedley, C. Fraser, S. Riley, L. J. Abu-Raddad, L. Ho, T. Thach, P. Chau, K. Chan, T. Lam, L. Tse, T. Tsang, S. Liu, J. H. Kong, E. M. Lau, N. M. Ferguson, and R. M. Anderson. Epidemiological determinants of spread of causal agent of severe acute respiratory syndrome in hong kong. *The Lancet*, 361(9371):1761–1766, 2003.

[19] S. Eubank, H. Guclu, V. S. A. Kumar, M. V. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429(6988):180–184, 2004.

[20] N. M. Ferguson, D. A. T. Cummings, S. Cauchemez, C. Fraser, S. Riley, A. Meeyai, S. Iamsirithawor, and D. S. Burke. Strategies for containing an emerging influenza pandemic in Southeast Asia. *Nature*, 437(7056):204–214, 2005.

[21] M. Ferrari, S. Bansal, L. A. Meyers, and O. Bjornstad. Network frailty and the geometry of herd immunity. *Proceedings of the Royal Society B*, 273(1602):2743–2748, 2006.

[22] K. Ford, W. Sohn, and J. Lepkowski. American adolescents: sexual mixing paterns, bridge partners, and concurrency. *Sexually Transmitted Diseases*, 29(1):13–19, 2002.

[23] C. Fraser, C. A. Donnelly, S. Cauchemez, W. P. Hanage, M. D. Van Kerkhove, T. D. Hollingsworth, J. Griffin, R. F. Baggaley, H. E. Jenkins, E. J. Lyons, T. Jombart, W. R. Hinsley, N. C. Grassly, F. Balloux, A. C. Ghani, N. M. Ferguson, A. Rambaut, O. G. Pybus, H. Lopez-Gatell, C. M. Alpuche-Aranda, I. B. Chapela, E. P. Zavala, D. M. E. Guevara, F. Checchi, E. Garcia, S. Hugonnet, C. Roth, and The WHO Rapid Pandemic Assessment Collaboration. Pandemic potential of a strain of influenza A (H1N1): Early findings. *Science*, 324(5934):1557–1561, 2009.

[24] T. C. Germann, K. Kadau, I. M. Longini, and C. A. Mackan. Mitigation strategies for pandemic influenza in the United States. *Proceedings of the National Academy of Sciences*, 103(15):5935–5940, 2006.

[25] S. Goll. Design and Implementation of the Disease Control System DiCon. Master's thesis, The University of Texas at Austin, Austin, Texas, December 2009.

[26] P. Grassberger. On the critical behavior of the general epidemic process and dynamical percolation. *Mathematical Biosciences*, 63(2):157–172, 1983.

[27] G. Grimmett. *Percolation*. Springer, Berlin, 1999.

[28] Health Protection Agency, Health Protection Scotland, National Public Health Service for Wales, and HPA Northern Ireland Swine Influenza Investigation teams. Epidemiology of new influenza A (H1N1) virus infection, United Kingdom, April – June 2009. *Eurosurveillance*, 14(22):pii=19232, 2009.

[29] M. J. Keeling and P. Rohani. *Modeling Infectious Diseases in Humans and Animals*. Princeton University Press, Princeton, New Jersey, 2007.

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

[31] J. Kleinberg and S. Lawrence. The structure of the web. *Science*, 294(5548):1849–1850, 2001.

[32] L. Kocsis and C Szepesvári. *Bandit Based Monte-Carlo Planning*, pages 282–293. Springer, New York, New York, 2006.

[33] D. J. Lanska. The mad cow problem in the UK: risk perceptions, risk management, and health policy development. *Journal of Public Health Policy*, 19(2):160–183, 1998.

[34] F. Liljeros, C. R. Edling, L. A. N. Amaral, and H. E. Stanley. The web of human sexual contacts. *Nature*, 411(6840):907–908, 2001.

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

[36] E. F. Long and M. L. Brandeau. OR's next top model: Decision models for infectious disease control. *Tutorials in Operations Research*, pages 123–138, 2009.

[37] I. M. Longini, M. E. Halloran, A. Nizam, and Y. Yang. Containing pandemic influenza with antiviral agents. *American Journal of Epidemiology*, 159(7):623–633, 2004.

[38] S. Mannor and J. Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5(Jun):623–648, 2004.

[39] L. A. Meyers. Contact network epidemiology: Bond percolation applied to infectious disease prediction and control. *Bulletin of the American Mathematical Society*, 44(1):63–86, 2007.

[40] L. A. Meyers, M. E. J. Newman, M. Martin, and S. Schrag. Applying network theory to epidemics: control measures for mycoplasma pneumoniae outbreaks. *Emerging Infectious Diseases*, 9(2):204–210, 2003.

[41] L. A. Meyers, M. E. J. Newman, and B. Pourbohloul. Predicting epidemics on directed contact networks. *Journal of Theoretical Biology*, 240(3):400–418, 2006.

[42] L. A. Meyers, B. Pourbohloul, M. E. J. Newman, D. M. Skowronski, and R. C. Brunham. Network theory and sars: predicting outbreak diversity. *Journal of Theoretical Biology*, 232(1):71–81, 2005.

[43] J. Mossong, N. Hens, M. Jit, P. Beutels, K. Auranen, R. Mikolajczyk, M. Massari, S. Salmaso, G. S. Tomba, J. Wallinga, J. Heijne, M. Sadowska-Todys, M. Rosinska, and W. J. Edmunds. Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLoS Medicine*, 5(3):e74, 2008.

[44] M. E. J. Newman. Assortative mixing in networks. *Physical Review Letters*, 89(20):208701, 2002.

[45] M. E. J. Newman. Spread of epidemic disease on networks. *Physical Review E*, 66(1):016128, 2002.

[46] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Physical Review Letters*, 86(14):3200–3203, 2001.

[47] B. Pourbohloul, A. Ahued, B. Davoudi, R. Meza, L. A. Meyers, D. M. Skowronski, I. Villasen, F. Galva, P. Cravioto, D. J. D. Earn, J. Dushoff, D. Fisman, W. J. Edmunds, N. Hupert, S. V. Scarpino, J. Trujillo, M. Lutzow, J. Morales, A. Contreras, Chávez C., D. M. Patrick, and R. C. Brunhama. Initial human transmission dynamics of the pandemic (H1N1) 2009 virus in North America. *Influenza and Other Respiratory Viruses*, 3(5):215–222, 2009.

[48] B. Pourbohloul, L. A. Meyers, D. M. Skowronski, M. Krajden, D. M. Patrick, and R. C. Brunham. Modeling control strategies of respiratory pathogens. *Emerging Infectious Diseases*, 11(8):1249–1256, 2005.

[49] J. M. Read, K. T. D. Eames, and W. J. Edmunds. Dynamic social networks and the implications for the spread of infectious disease. *Journal of the Royal Society Interface*, 5(26):1001–1007, 2008.

[50] S Riley, C. Fraser, C. A. Donnelly, A. C. Ghani, L. J. Abu-Raddad, A. J. Hedley, G. M. Leung, L. Ho, Lam T., T. Q. Thach, P. Chau, K. Chan, S. Lo, P. Leung, T Tsang, W. Ho, K. Lee, E. M. C. Lau, N. M. Ferguson, and Anderson R. M. Transmission dynamics of the etiological agent of SARS in Hong Kong: Impact of public health interventions. *Science*, 300(5627):1961–1966, 2003.

[51] Å. Svensson. A note on generation times in epidemic models. *Mathematical Biosciences*, 208(1):300–311, 2007.

[52] Texas Department of State Health Services. *Planning Guidelines for Nonpharmaceutical Interventions*. Texas Department of State Health Services, Austin, Texas, 2007.

[53] The Scottish Government. Update on A (H1N1) virus. http://www.scotland.gov.uk/News/Releases/2009/06/01153403, June 2009. Accessed on 08/24/2010.

[54] U.S. Department of Health and Human Services. HHS unveils two new efforts to advance pandemic flu preparedness. http://www.hhs.gov/news/press/2007pres/20070201.html, February 2007. Accessed on 08/24/2010.

[55] E. Volz. SIR dynamics in random networks with heterogeneous connectivity. *Journal of Mathematical Biology*, 56(3):293–310, 2007.

[56] E. Volz and L. A. Meyers. Susceptible-infected-recovered epidemics in dynamic contact networks. *Proceeding of the Royal Society B*, 274(1628):2925–2933, 2007.

[57] E. Volz and L. A. Meyers. Epidemic thresholds in dynamic contact networks. *Journal of the Royal Society Interface*, 6(32):233–241, 2009.

[58] J. Wallinga and P. Teunis. Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *American Journal of Epidemiology*, 160(6):509–516, 2004.

[59] C. Watts and R. May. The influence of concurrent partnerships on the dynamics of HIV/AIDS. *Mathematical Biosciences*, 108(1):89–104, 1992.

[60] D. J. Watts. *Small Worlds: The Dynamics of Networks Between Order and Randomness.* Princeton University Press, 1999.

[61] D. J. Watts. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences of the United States of America*, 99(9):5766–5771, 2002.

[62] L. F. White and M Pagano. A likelihood-based method for real-time estimation of the serial interval and reproductive number of an epidemic. *Statistics in Medicine*, 27(16):2999–3016, 2008.

[63] World Health Organization. Update 49 - SARS case fatality ratio, incubation period. http://www.who.int/csr/sarsarchive/2003_05_07a/en/, May 2003. Accessed on 08/24/2010.

[64] World Health Organization. Summary of probable SARS cases with onset of illness from 1 November 2002 to 31 July 2003. http://www.who.int/csr/sars/country/table2004_04_21/en/index.html, May 2004. Accessed on 08/24/2010.

[65] Y. Yang, J. D. Sugimoto, E. Halloran, N. E. Basta, D. L. Chao, L. Matrajt, G. Potter, E. Kenah, and I. M. Longini. The transmissibility and control of pandemic influenza A (H1N1) virus. *Science*, 326(5953):729–733, 2009.

A   Geometric and Exponential Random Variables

Consider the following sequence of geometric random variables. The variable X_1 is the standard geometric random variable, with success probability β and takes values in $\{1, 2, \ldots\}$. The variable X_2 has success probability $\frac{\beta}{2}$ and takes values in $\{\frac{1}{2}, \frac{2}{2}, \ldots\}$. Similarly, the variable X_k has success probability $\frac{\beta}{k}$ and takes values in $\{\frac{1}{k}, \frac{2}{k}, \ldots\}$ and so forth.

Consider

$$\lim_{k \to \infty} \Pr[X_k \leq y] = \lim_{k \to \infty} \sum_{i=0}^{\frac{i}{k} \leq y - \frac{1}{k}} (1 - \frac{\beta}{k})^i \cdot \frac{\beta}{k}.$$

For the upper limit of the summation on the right-hand-side, i can go as high as $k \cdot y - 1$. Substituting, we have

$$\begin{aligned}
\lim_{k \to \infty} \Pr[X_k \leq y] &= \lim_{k \to \infty} \sum_{i=0}^{k \cdot y - 1} (1 - \frac{\beta}{k})^i \cdot \frac{\beta}{k} \\
&= \lim_{k \to \infty} \frac{1 - (1 - \frac{\beta}{k})^{k \cdot y}}{1 - (1 - \frac{\beta}{k})} \cdot \frac{\beta}{k} \\
&= \lim_{k \to \infty} 1 - (1 - \frac{\beta}{k})^{k \cdot y} \\
&= 1 - e^{-\beta \cdot y}.
\end{aligned}$$

So, the sequence of geometric random variables converges in distribution to an exponential random variable with expectation $\frac{1}{\beta}$.

The idea here is that in the sequence of geometrics, we decreased the spacing between their values and the success probability with the same scale, keeping the expected number of successes in a unit of time the same.

B   Generation Time of an SIR Model

As a small example of the complex and model-dependent methods of parameter estimation, let us express the generation time of the SIR model discussed in Section 2. In an SIR model, the infectious period of an individual is an exponentially distributed random variable, with expectation $\frac{1}{\gamma}$, denoted as $\text{Exp}(\gamma)$. Consider an index case at the beginning of the epidemic, when S is approximately 1. The infected individual produces new infected individuals with a waiting-time between new infections that is an exponential random variable with expectation $\frac{1}{\beta}$, denoted as $\text{Exp}(\beta)$. If the index case is infected at time 0, and given that at least one new infection was created, what is the average time of infection for the secondary cases?

Let $Y \sim \text{Exp}(\gamma)$ and $Z \sim \text{Exp}(\beta)$. The variable Y captures the infectious period of the index case, and Z captures the waiting time between generation of consecutive secondary cases. We can view the creation of new infected cases as a repeated "race" between Y and Z. If Z wins the race, by being smaller than Y, then a new infected case is created and a new race is started. If Y wins the race, then the index case recovers, the repeated races stop, and no new infected cases can be created.

Define X_1 to be the infection time of the first newly created case. Define X_2 to be the difference in the infection times between the first and second newly infected cases, so that the second newly infected case occurs at time $X_1 + X_2$. Similarly, define X_3 to be the difference in infection times between the second and third, and so forth for X_i for $i = 1 \dots$. Further, let P_i be the probability that i new cases are created, for $i = 1 \dots$. With these definitions, the generation time is equal to

$$E\left[\sum_{i=1}^{\infty} P_i \cdot \frac{\sum_{j=1}^{i} \sum_{k=1}^{j} X_k}{i}\right] \tag{15}$$

where with the innermost sum, we calculate the infection time of the jth new case; with the second innermost sum, we calculate the average over all cases created; and finally, with the outermost sum, we calculate the expectation over the total number of newly created cases.

To analyze Expression (15), first, recall that if $Y \sim \text{Exp}(\gamma)$ and $Z \sim \text{Exp}(\beta)$, then $\min(Y, Z) \sim \text{Exp}(\gamma + \beta)$. Also recall that an exponential random variable has the "lack of memory" property. In specific $(Y - y \mid Y \geq y) \sim \text{Exp}(\gamma)$, for all non-negative constants y.

Now, suppose that exactly one new case is created, then the new case created at time $X_1 = \min(Y, Z)$, so $X_1 \sim \text{Exp}(\gamma + \beta)$. If exactly two new cases are created, a new race between Y and Z is started after time X_1, because of Y's lack of memory property. Because of the new race, we also have $X_2 = \min(Y, Z)$, and $X_2 \sim \text{Exp}(\gamma + \beta)$. In fact, all X_i have the same distribution, $\text{Exp}(\gamma + \beta)$. This allows us to greatly simplify Expression (15) to

$$E\left[\sum_{i=1}^{\infty} P_i \cdot \frac{\sum_{j=1}^{i} \sum_{k=1}^{j} X_k}{i}\right]$$

$$= \sum_{i=1}^{\infty} P_i \cdot E[X_1] \cdot \frac{\sum_{j=1}^{i} j}{i}$$

$$= \sum_{i=1}^{\infty} P_i \cdot E[X_1] \cdot \frac{\frac{i \cdot (i+1)}{2}}{i}$$

$$= \frac{1}{\gamma + \beta} \cdot \sum_{i=1}^{\infty} P_i \cdot \frac{(i+1)}{2} \tag{16}$$

We can complete our analysis of the generation time by noting that the number of new infected individuals, which is the number of races between Y and Z in our analogy, is geometrically distributed. The variable Y has a chance of winning a given race equal to $\frac{\beta}{\gamma+\beta}$, which can be derived with the appropriate integral. Whenever Y wins, the races stop, giving our geometric distribution. Since we are given that at least one secondary case has occurred, we can use a normalized geometric distribution to simply Expression (16) to

$$\frac{1}{\gamma+\beta} \cdot \sum_{i=1}^{\infty} P_i \cdot \frac{(i+1)}{2}$$

$$= \frac{1}{\gamma+\beta} \cdot \sum_{i=1}^{\infty} \frac{(\frac{\gamma}{\gamma+\beta})^i \cdot \frac{\beta}{\gamma+\beta}}{1 - \frac{\beta}{\gamma+\beta}} \cdot \frac{(i+1)}{2}$$

$$= \frac{\beta}{2\gamma(\gamma+\beta)} \cdot \sum_{i=1}^{\infty} (\frac{\gamma}{\gamma+\beta})^i (i+1) \tag{17}$$

Finally, we can use the fact that $\sum_{i=1}^{\infty} r^i \cdot (i+1) = \frac{r}{1-r} + \frac{r}{(1-r)^2}$ for $r < 1$ to simplify Expression (17) to derive our final expression of the generation time of an SIR model

$$\frac{1}{2(\gamma+\beta)} + \frac{1}{2\gamma}$$

The correctness of the expression can be easily verified through simulation.

This small example demonstrates the complex reasoning required for proper parameterization of epidemic models. Simple quantities like generation time or serial number are what can be estimated for real disease cases. To parameterize epidemic models, one must connect these real estimates to the model parameters. This complexity combined with the importance of using the correct parameters is perhaps one of the reasons that a large fraction of the literature on infectious diseases is concentrated on parameter estimation.
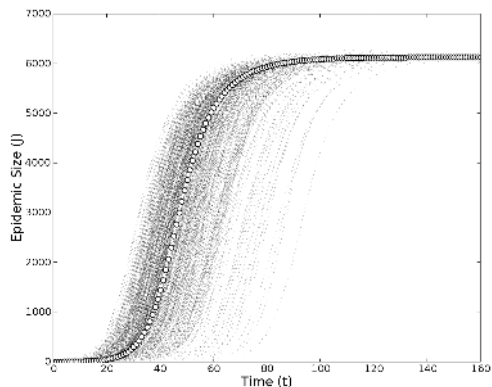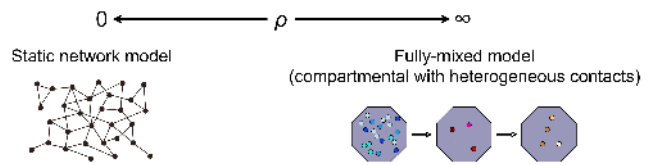
(a)   (b)

Figure 7: Disease spread through a dynamic contact network. (a) We compare our mathematical model to stochastic simulations. The mathematical predictions (circles) fall right in the center of mass of 1000 stochastic simulations (dotted lines) of disease transmission through a dynamic network. This assumes a Poisson network with mean degree 1.5, transmissability $r = 0.2$, recovery $\mu = 0.1$, and mixing rate $\rho = 0.25$. (b) As the mixing rate increases from zero to infinity, the model smoothly interpolates between static network models and compartmental models.
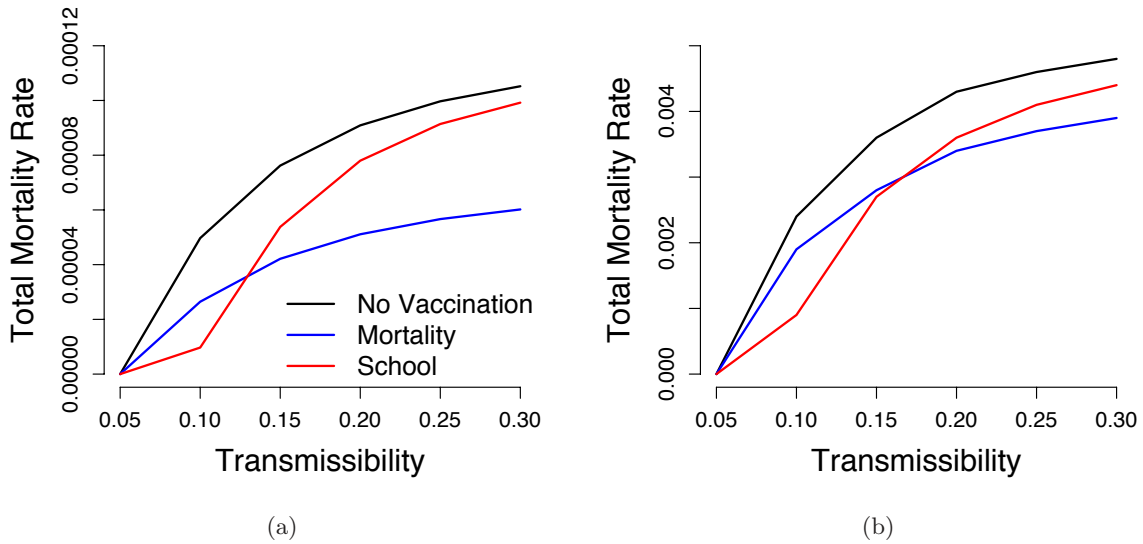
Figure 8: Prioritizing flu vaccines under limited supplies. This study assumes that there are enough vaccines to cover 13% of the population (as occurred during the 2004-2005 influenza vaccination shortage) and compares the efficacies of prioritizing school children (red) versus groups at high risk for mortality from flu (blue). We use a contact network model based on detailed sociological and demographic information for the city of Vancouver, British Columbia; and model vaccination by removing effectively vaccinated nodes and their edges from the network (using published estimates for age-specific influenza vaccine efficacy). The impacts of such removals are predicted using the bond percolation methods described above. (a) Seasonal flu model in which high risk groups include infants and elderly. (b) Pandemic flu model (based on the 1918-1919 Spanish Flu Pandemic) in which adults have the highest mortality rates followed by infants. For a full description, see Bansal et al. [9] The x-axes give T, the average transmissability of influenza. Estimates for for T vary across the interval from 0.10 to 0.30 for both seasonal and pandemic flu. In both cases, prioritizing school children is predicted to cause a greater reduction in mortality than prioritizing high-risk groups for mildly contagious flus (low T), while the reverse is true for more highly contagious strains (high T). The transition between these two outcomes occurs at a slightly higher transmissability for pandemic flu than for seasonal flu.
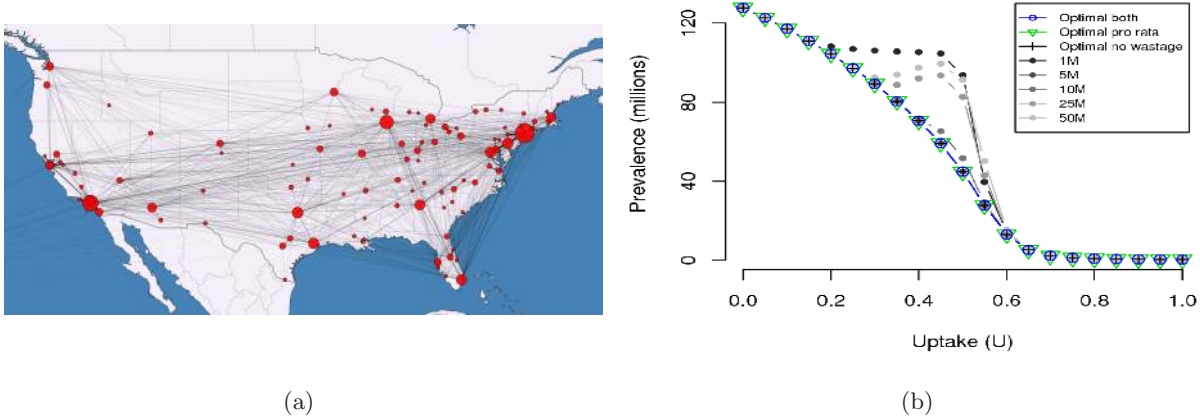
Figure 9: Optimizing the distribution of antiviral medications from the US Strategic National Stockpile (SNS). (a) The network model used to stochastically simulate influenza progression throughout the 100 most populated U.S. cities. Transmission within cities is modeled using compartmental models and transmission among cities occurs via stochastic movement of infected travelers. The model is parameterized with data from the U.S. Census Bureau, the Bureau of Transportation Statistics, and early estimates of H1N1/09 parameters [47]. Node (circle) size is proportional to city population and edge thickness is proportional to the number of daily travelers. (b) The performance of several control strategies in delaying an influenza epidemic. The x-axis gives rates of *uptake*, the fraction of individuals who seek antivirals within the first 24 hours of symptoms. On the vertical axis is the cumulative number of infected cases within the first 12 months of the initial outbreak. The model assumes that following the distribution of antivirals to cities, they disappear through misuse or loss, with a half life of two months. Three optimized strategies for releasing 50M courses from the national antiviral stockpile are presented (from top to bottom): the optimized strategy when antivirals are allowed to be released both proportional to population size and proportional to disease prevalence, the optimized strategy when releases are exclusively proportional to city population sizes, and the optimized strategy under an ideal situation assuming that there is no misuse (infinite half life). In addition, several simple policies are presented ranging from a monthly release of 1 million antiviral courses for 12 months to a single release of the entire 50M antiviral courses available in the stockpile. The results suggest that: 1) releases proportional to prevalence are unnecessary, since the performance is the same with or without this option, 2) careful release can overcome misuse or loss, since the best release schedules under misuse perform as well as the idealized scenario of no misuse, and 3) the simple strategy of releasing 5M courses monthly performs well for the H1N1/09 disease parameters. The same performance not necessarily occur for other strains of influenza with different characteristics.