

Network theory and SARS: predicting outbreak diversity

Lauren Ancel Meyers^{a,b,*}, Babak Pourbohloul^{c,1,2}, M.E.J. Newman^{b,d},
Danuta M. Skowronski^{c,2}, Robert C. Brunham^{c,2}

^aSection of Integrative Biology and Institute for Cellular and Molecular Biology, University of Texas at Austin, 1 University Station C0930, Austin, TX 78712, USA

^bSanta Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

^cUniversity of British Columbia Centre for Disease Control, 655 West 12th Avenue, Vancouver, British Columbia, Canada V5Z 4R4

^dCenter for the Study of Complex Systems, University of Michigan, Randall Laboratory, 500 E. University Ave., Ann Arbor, MI 48109-1120, USA

Received 3 May 2004; received in revised form 6 July 2004; accepted 22 July 2004

Available online 23 September 2004

Abstract

Many infectious diseases spread through populations via the networks formed by physical contacts among individuals. The patterns of these contacts tend to be highly heterogeneous. Traditional “compartmental” modeling in epidemiology, however, assumes that population groups are fully mixed, that is, every individual has an equal chance of spreading the disease to every other. Applications of compartmental models to Severe Acute Respiratory Syndrome (SARS) resulted in estimates of the fundamental quantity called the basic reproductive number R_0 —the number of new cases of SARS resulting from a single initial case—above one, implying that, without public health intervention, most outbreaks should spark large-scale epidemics. Here we compare these predictions to the early epidemiology of SARS. We apply the methods of contact network epidemiology to illustrate that for a single value of R_0 , any two outbreaks, even in the same setting, may have very different epidemiological outcomes. We offer quantitative insight into the heterogeneity of SARS outbreaks worldwide, and illustrate the utility of this approach for assessing public health strategies.

© 2004 Elsevier Ltd. All rights reserved.

Keywords: SARS; Epidemiology; Intervention; Contact network; Mathematical model

1. Introduction

More than two years since the first case of severe acute respiratory syndrome (SARS), a respiratory illness caused by a novel coronavirus, occurred in Guangdong province of China (November, 2002) and more than 18

months since the syndrome was first recognized outside of Asia (in Canada on March 13, 2003), its pattern of spread remains an enigma to public health officials and epidemiologists (Cyranoski and Abbott, 2003; Drosten et al., 2003; Ksiazek et al., 2003; Marra et al., 2003; Peiris et al., 2003; World Health Organization, 2003). Mathematical epidemiologists originally estimated the average number of secondary cases emanating from one primary case in a susceptible population (R_0) to be in the range of 2.2 to 3.6 for this virus—an estimate well above one, approximating that of a new subtype of influenza (Hethcote, 2000; Lipsitch, 2003; Riley et al., 2003).

Despite this estimate and near-universal susceptibility, SARS has not emerged as a global pandemic. Instead, initial seeding was followed by intense but

*Corresponding author. Section of Integrative Biology, University of Texas, 1 University Station #C0930, Austin, TX 78712, USA. Tel.: +1-512-471-4950; fax: +1-512-471-3878.

E-mail address: laurenmeyers@mail.utexas.edu (L.A. Meyers).

¹These authors contributed equally to the conception and delivery of this work.

²Department of Health Care and Epidemiology, Faculty of Medicine, University of British Columbia, 5804 Fairview Avenue, Vancouver, BC, Canada V6T 1Z3.

tightly circumscribed activity in some locales with only scant activity in others. In Canada, for instance, Toronto, Ontario and Vancouver, British Columbia were first affected nearly simultaneously in March 2003. By June 3, 2003, Toronto had experienced more than 209 probable cases and Vancouver had experienced only four probable cases. No other province of Canada reported any probable cases (Health Canada, 2003). The United States with a population more than fifty-fold greater than Toronto reported 69 probable cases—67 imported and only two from secondary spread (Centers for Disease Control and Prevention, 2003).

The discrepancy between the estimates of R_0 and the observed epidemiology might stem from early and effective intervention since R_t , the reproductive ratio of a disease at time t , will decrease with the implementation of successful infection control measures. Yet, even during the three and a half months of SARS spread in China between its initial appearance and the broad implementation of public health measures, case counts were much less than expected from such values of R_0 (Xu et al., 2004). By definition, the total number of expected cases of a disease goes up by a factor of R_0 for every generation of infection, a generation being the mean time between an individual becoming infected and their infecting others. Based on recorded dates of the first symptoms for 124 pairs of infections in Singapore and Toronto (Leo et al., 2003; Poutanen et al., 2003), we estimate the average generation time (γ) for SARS to be 9.7 ± 0.3 days. (This estimate clearly depends on the accuracy of the reported data.) Roughly, the cumulative number of SARS cases over D days should be $\sum_{i=0}^{D/\gamma} (R_0)^i = (1 - R_0^{D/\gamma+1}) / (1 - R_0)$ (This is capped by total population size and does not consider the reduction in R_t once a substantial proportion of the population is infected). Thus for R_0 ranging between 2.2 and 3.6, this equation predicts that in the first 120 days of transmission in China, there should have been between approximately 30,000 and 10 million cases. In fact only 782 cases were reported during the initial three months (World Health Organization, 2003), which, using this simple calculation, suggests that R_0 should be much lower and closer to 1.6. A subsequent estimate based on data from the Hong Kong and Singapore outbreaks brings R_0 down to 1.2, which, by the above formula, predicts approximately 50 cases during the first 120 days of transmission (Chowell et al., 2003). While this number agrees nicely with the case counts observed in Hong Kong and Singapore (Chowell et al., 2003a, b) it is an order of magnitude lower than that reported for China.

Why do epidemiologists derive such varied estimates of R_0 and why were the initial estimates so high in comparison to the observed epidemiology in China? The basic premise of fully mixed epidemiological models—that all individuals in a group are equally likely to

become infected (or infect others)—often does not hold and therefore may lead to spurious estimates or estimates that cannot justifiably be extrapolated from the specific setting in which they were measured to the broader community context. Early SARS estimates were based largely on transmission data from closed settings like hospitals and crowded apartment buildings, where there are unusually high rates of contact between individuals (Lipsitch, 2003; Riley et al., 2003). In fact, hospital transmission accounts for 50% of the value of R_0 described in Riley et al. (2003). If the contact patterns in these settings were highly heterogeneous, then the estimates for R_0 may be inaccurate. Even if the estimates for R_0 were indeed appropriate for these specific settings, they probably should not be extrapolated to the population at large. Contact rates may be considerably lower outside hospitals and crowded apartment buildings and thus so may be the general value of R_0 for SARS (Yu et al., 2004). Such disparity may account for the discrepancy between the estimates and the slower progress of the outbreak in China. In fact, further studies suggest that the unusually large cluster of infected cases in Amoy Gardens complex in Hong Kong was due to exposure to the virus-laden aerosol plume originating from one of the buildings in that area and not from direct person-to-person contact (Yu et al., 2004). A recent analysis of the impact of SARS isolation interventions in Hong Kong, Singapore and Toronto emphasizes the importance of viewing R_0 as a distribution of possible values where the mean and median may vary depending on the setting in which the disease is spreading (Chowell et al., 2004).

SARS, like many other infectious diseases exhibits great heterogeneity in transmission efficiency with certain individuals appearing to be responsible for large proportion of transmission events (Booth et al., 2003; Donnelly et al., 2003; Leo et al., 2003). These individuals may be “superspreaders” with unusually large numbers of contacts or “supershedders” who are unusually effective at excreting the virus into the environment they share with others. In contrast to the fully mixed assumption of standard compartmental models, the contact patterns in a community may be quite diverse. There is an enormous difference between a situation in which all individuals share typical contact patterns and one in which most infected individuals pass the disease on to only one or even zero others, but a small number pass it onto dozens or even hundreds—the mean value of R_0 can be the same in both cases, while the epidemiological outcomes are vastly different.

Public health control measures for communicable diseases, including contact tracing, isolation, quarantine, and ring vaccination (Greenhalgh, 1986; Müller et al., 2000) have historically been predicated on sociological considerations. For example, the strategies for international SARS control were founded upon the

largely intuitive idea of early interruption of critical social contacts (World Health Organization, 2003). There are several mathematical approaches that likewise consider sociological factors. To name a few, compartmental models that break populations into many demographic groups capture greater behavioral heterogeneity (Hethcote and Yorke, 1984); stochastic approaches including stochastic compartmental models (Bailey, 1975), branching process models (Becker, 1977; Farrington et al., 2003), dyad models (Keeling et al., 1997; Ferguson and Garnett, 2000), and Reed-Frost chain-binomial models (Lefevre and Picard, 1989) allow more exact predictions of the size and probability of epidemics; and “individual-based modeling”, a primarily computational approach based on following the contact and infection histories of simulated individuals, yields detailed statistical predictions about disease outcomes. Most individual-based models assume predefined simple contact patterns such as regular lattices (Durrett, 1999; Kleczkowski and Grenfell, 1999; Ritton and O’neill, 2002; Sander et al., 2002), although there have been recent efforts to consider more realistic contact patterns (Van der Ploug et al., 1998; Chowell et al., 2003a, b; Eubank et al., 2004).

A recent addition to this toolkit is *contact network epidemiology*, an analytical framework that explicitly captures the diverse interactions that underlie the spread of diseases (Longini 1988; Sattenspiel and Simon, 1988; Morris, 1995; Ball et al., 1997; Diekmann et al., 1998; Lloyd and May, 2001; Newman, 2002; Keeling et al., 2003; Meyers et al., 2003). The first step in contact network epidemiology is to construct networks based on information about real-life contacts between individuals. One can then analyse these networks to determine their crucial topological features and apply analytical methods to make epidemiological predictions and intervention recommendations. The two primary advantages of this approach are that (i) it makes no a priori assumptions about the global network structure and (ii) the mathematical analysis allows one to bypass extensive simulation.

Here, we extend the mathematics of contact network epidemiology to make more detailed public health predictions, to demonstrate the importance of an accurate model of contact patterns in a community, and to provide new insight into the observed epidemiology of SARS.

2. Contact networks

Contact network models attempt to characterize every interpersonal contact that can potentially lead to disease transmission in a community. These contacts may take place within households, schools, workplaces, hospitals, etc. (Fig. 1A). Each person in a community is

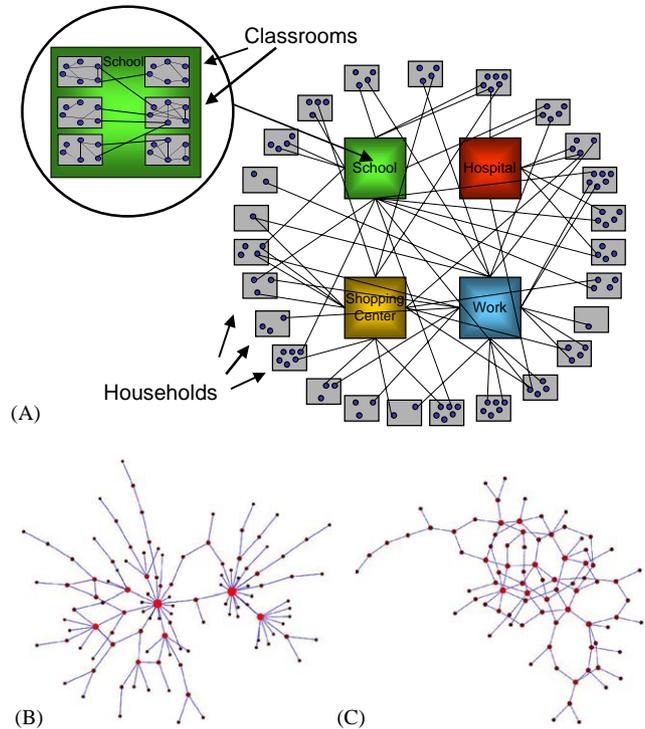


Fig. 1. Schemata of: (A) urban, (B) power law, and (C) Poisson networks. Dots represent individuals and lines between dots represent contacts between individuals that could potentially lead to disease transmission.

represented as a *vertex* in the network and each contact between two people is represented as an *edge* (line) connecting their vertices. The number of edges emanating from a vertex, that is, the number of contacts a person has, is called the *degree* of the vertex. The distribution of the numbers of contacts—the degree distribution—is a fundamental quantity in network theory.

In the studies described here, we start by generating a plausible contact network for an urban setting using computer simulations. The simulations are based on data for the city of Vancouver, British Columbia. We choose $N=1000$ households at random from the Vancouver household size distribution (Statistics Canada, 2001), which yields approximately 2600 people. Household members are given ages according to the measured Vancouver age distribution (BC Stats, 2003), and, based on age, are then assigned to schools according to school and class size distributions (Vancouver School Board, 2002), to occupations according to (un)employment data (BC Stats, 2002), to hospitals as patients and caregivers according to hospital employment and bed data (Centre for Health Services and Policy Research, 2002), and to other public places. Within each location we create random connections between individuals with probabilities of 1 for households, 0.3 for hospitals and schools, 0.03 for workplaces, and 0.003 for other public places. Admittedly, these

within-location parameters are largely based on intuition, and future versions of this model will include estimates based on data. We found that the structure of these simulated networks is robust to small changes in these values. Each school or hospital is sub-divided into classrooms or wards. Pairs of students or patients within these sub-units were connected with higher probability than pairs associated with different sub-units. Teachers are assigned to classrooms and connected stochastically to appropriate students. Caregivers are assigned wards and then connected to appropriate patients. There are also low probability neighborhood contacts between individuals from different households³.

This network offers a high degree of realism but is quite complex. We therefore use two simpler networks to provide additional insight. One is a random network with a Poisson degree distribution in which individuals connect to others independently and uniformly at random (Fig. 1B). Neither the simulated nor the Poisson network, however, contains a significant fraction of superspreaders. Therefore we also study a network with a (truncated) power-law degree distribution (Fig. 1C), a form much discussed in recent work on network epidemiology (Barabasi and Albert, 1999; Pastor-Satorras and Vespignani, 2001). This network has a “heavy tail” of superspreaders (Fig. 2) and, as we will see, these individuals can have a profound effect on outbreak patterns despite being few in number. Network theorists often refer to such networks as scale-free because of the absence of a typical degree in the network.

We define the transmissibility of a disease, T , to be the average probability that an infectious individual will transmit the disease to a susceptible individual with whom they have contact. T summarizes core aspects of disease transmission including the rate at which contacts take place between individuals, the likelihood that a contact will lead to transmission, the duration of the infectious period, and the susceptibility of individuals to SARS. The epidemic threshold, which, in an uncorrelated network, is given by

$$T_c = \frac{\langle k \rangle}{\langle k^2 \rangle - \langle k \rangle},$$

where $\langle k \rangle$ and $\langle k^2 \rangle$ are the mean degree and mean square degree of the network, respectively, is the minimum transmissibility (T) required for an outbreak to become a large-scale epidemic.

We choose the parameters of the Poisson and power law networks so that all three networks share the same epidemic threshold T_c . Let p_k denote the probability that a randomly selected individual in a network has degree k . Then, the Poisson network is given by $p_k =$

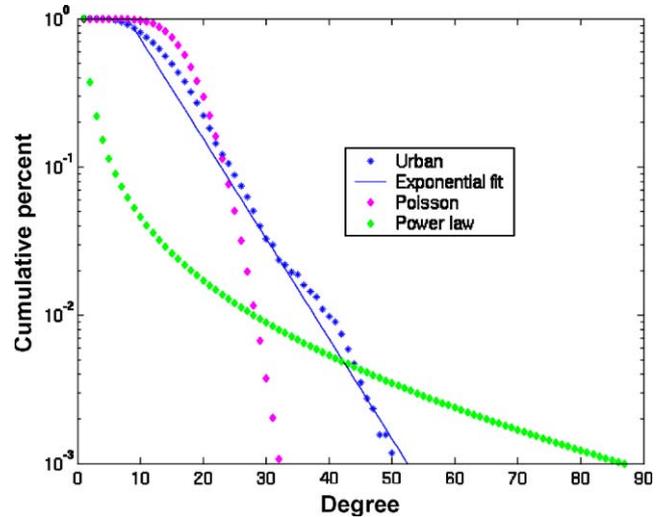


Fig. 2. Cumulative degree distributions for simulated urban, Poisson, and power law networks. As described in the text, these share the same epidemic threshold (T_c). Each line gives the probability that a randomly chosen individual (vertex) will have at least the number of contacts (degree) indicated on the x-axis. The degree distribution for the urban network is nearly exponential for degrees greater than ten.

$(\mu^k/k!) \exp(-\mu)$ with mean contact number $\mu = 19.6$; and the power law network is given by

$$p_k = \begin{cases} 0 & \text{for } k = 0, \\ Ck^{-\alpha} \exp(-\frac{k}{\kappa}) & \text{for } k > 0, \end{cases}$$

with distribution parameters $\kappa = 94.2$ and $\alpha = 2$. Here we fixed α and solved for κ . The results described below are qualitatively similar for a large range of values of α . Truncation of the power law distribution raises the epidemic threshold of the network to values comparable to those found for urban networks.

To generate the two idealized networks, we begin with a specified number of vertices and choose degrees for these vertices at random from the desired degree distribution. Then we connect random pairs of vertices, until the chosen degrees are exhausted. This often yields imperfect graphs with loops connecting vertices to themselves or redundant edges that connect the same two vertices more than once. We remove these imperfections using an algorithm suggested by Maslov et al. (2001) in which we select at random two edges connecting, for example, vertex pairs AB and CD, and swap them so that they now connect AC and BD, unless this would create a new loop or double edge, in which case we do nothing. This process occasionally eliminates loops and repeated edges and by repeating it a sufficiently large number of times (depending on the network size) we can produce a network with none at all.

2.1. Epidemiological analysis

Given the degree distribution of a network, we can use tools based on percolation theory (Pastor-Satorras

³The authors will provide copies of the simulation software and a detailed list of parameter values to any interested readers upon request.

and Vespignani, 2001; Newman, 2002) to predict the fate of an outbreak of an infectious disease as a function of its transmissibility T (defined above). T is related to the traditional R_0 according to

$$R_0 = T(\langle k^2 \rangle / \langle k \rangle - 1),$$

where $\langle k \rangle$ and $\langle k^2 \rangle$ are the mean degree and mean square degree of the network, respectively. The critical transmissibility value T_c (also defined above), corresponds to $R_0 = 1$, above which a population is vulnerable to large scale epidemics (but is not guaranteed to experience an epidemic) and below which only small local outbreaks occur.

Network theory makes a technical distinction between outbreaks and epidemics. An outbreak is a causally connected cluster of cases that, by chance or because the transmission probability is low, dies out before spreading to the population at large. In an epidemic, on the other hand, the infection escapes the initial group of cases into the community at large and results in population-wide incidence of the disease. The crucial difference is that the size of an outbreak is determined by the spontaneous dying out of the infection, whereas the size of an epidemic is limited only by the size of the population through which it spreads.

To predict the fate of an outbreak, we use probability generating functions (pgf), quantities that describe probability distributions, and here, summarize useful information about network topology. The pgf for a degree distribution is

$$G_0(x) = \sum_{k=1}^{\infty} p_k x^k.$$

If we choose a random edge and follow it to the nearest vertex, then the pgf for the “excess degree”—the number of edges emanating from that vertex other than the one along which we arrived—is

$$G_1(x) = \frac{\sum_{k=1}^{\infty} k p_k x^{k-1}}{\sum_{k=1}^{\infty} k p_k}.$$

The average degree and average excess degree equal the derivatives of these expressions at $x=1$, that is,

$$\langle k \rangle = \sum_{k=1}^{\infty} k p_k$$

and

$$\langle k_e \rangle = \frac{\sum_{k=1}^{\infty} k(k-1)p_k}{\sum_{k=1}^{\infty} k p_k} = (\langle k^2 \rangle / \langle k \rangle - 1),$$

respectively.

The value of the epidemic threshold T_c , the predicted average size of an outbreak $\langle s \rangle$ and probability of an epidemic S were first derived in (Newman, 2002). By

nesting pgf’s for the number of new infections emanating from an infected vertex one can construct a pgf for the size of a outbreak, and hence derive the average size of a outbreak:

$$\langle s \rangle = 1 + \frac{T \langle k \rangle}{1 - T \langle k_e \rangle}.$$

This expression diverges when an outbreak becomes a large-scale epidemic. The epidemic threshold T_c (above) is the transmissibility value that marks this point.

The probability of a full-blown epidemic S is derived by first calculating the likelihood that a single infection will lead to only an outbreak instead of a full-blown epidemic, and then subtracting that value from one:

$$S = 1 - \sum_{k=1}^{\infty} p_k (1 + (u - 1)T)^k,$$

where u is the probability that the person at the end of an edge does not have the disease and is the solution to the equation

$$u = \frac{\sum_{k=1}^{\infty} k p_k (1 + (u - 1)T)^{k-1}}{\sum_{k=1}^{\infty} k p_k}.$$

We use numerical root finding methods to solve for u . S is also the expected proportion of the population that will be infected should an epidemic occur.

Here we extend these results to predict the fate of an outbreak based on its initial conditions—how many and what kinds of individuals are already infected. We refer to the first individual in a community to come down with an infectious disease as *patient zero*. The probability that a patient zero with degree k will start an epidemic, ε_k , is equal to the probability that transmission of the disease along at least one of the edges emanating from the original vertex will lead to an epidemic. For any one of its k edges, $1 - T$ is the probability that the disease does not get transmitted along the edge and Tu is the probability that even if disease is transmitted to the next vertex, it does not proceed into a full-blown epidemic. Thus

$$\varepsilon_k = 1 - (1 - T + Tu)^k.$$

The probability that an outbreak of size N will ignite an epidemic is $1 - \prod_{i=1}^N (1 - \varepsilon_{k_i})$ where k_i is the degree of individual i . This is just one minus the probability that none of the N infected individuals sparks an epidemic. If we know the number of current cases but not their contact patterns, then our best estimate for the probability of an epidemic is calculated similarly, with each of the $(1 - \varepsilon_{k_i})$ ’s replaced with the probability that a typical infected individual does not spark an epidemic. The number of edges through which a typical infected individual can start an epidemic is given by the excess degree pgf, and the probability that one of those edges will not give rise to an epidemic is $1 - T + Tu$. Thus the

probability that none of those edges will be a conduit to an epidemic is

$$\left(\frac{\sum_{k=1}^{\infty} k p_k (1-T+Tu)^{k-1}}{\sum_{k=1}^{\infty} k p_k} \right),$$

and the probability that an outbreak of size N sparks an epidemic is

$$1 - \left(\frac{\sum_{k=1}^{\infty} k p_k (1-T+Tu)^{k-1}}{\sum_{k=1}^{\infty} k p_k} \right)^N.$$

Finally, we derive an individual’s risk of infection during an epidemic as a function of his or her degree. The probability v_k that an individual with degree k will become infected during an epidemic is equal to one minus the probability that none of his or her k contacts will transmit the disease to him or her. The probability that a contact does not transmit the disease is equal to the probability that the contact was infected, but did not transmit the disease, $(1-u)(1-T)$, plus the probability that the contact was not infected in the first place, u . Thus, a randomly chosen vertex of degree k will become infected with probability

$$v_k = \varepsilon_k = 1 - (1 - T + Tu)^k.$$

2.2. Epidemic simulations

We have simulated disease spread to verify our analytical predictions. Beginning with a susceptible population and a single case or a small cluster of cases, we iteratively take each currently infected vertex, infect each of its susceptible neighbors with probability T and then change the status of the original vertex to recovered. This method of simulation does not capture the temporal progression of an epidemic, but just the overall number and distribution of infected individuals.

2.3. Predicting epidemics in the contact networks

We predicted the probability S that an outbreak with $R_0 > 1$ will lead to an epidemic for the three networks described in Figs. 1 and 2. S is often significantly less than one, and can be different for two networks with the same R_0 (Fig. 3). When S is well below one and R_0 is well above one, it is very likely that communities with similar contact patterns will have diverse experiences with the disease, some experiencing large epidemics and other experiencing only minor outbreaks. In particular, above the epidemic threshold, the probability of an epidemic in the power law network and the probability of an epidemic in either of the other networks diverge quickly. Outbreaks are consistently less likely to reach epidemic proportions in the power law network than in the others. The vertical lines in Fig. 3 correspond to

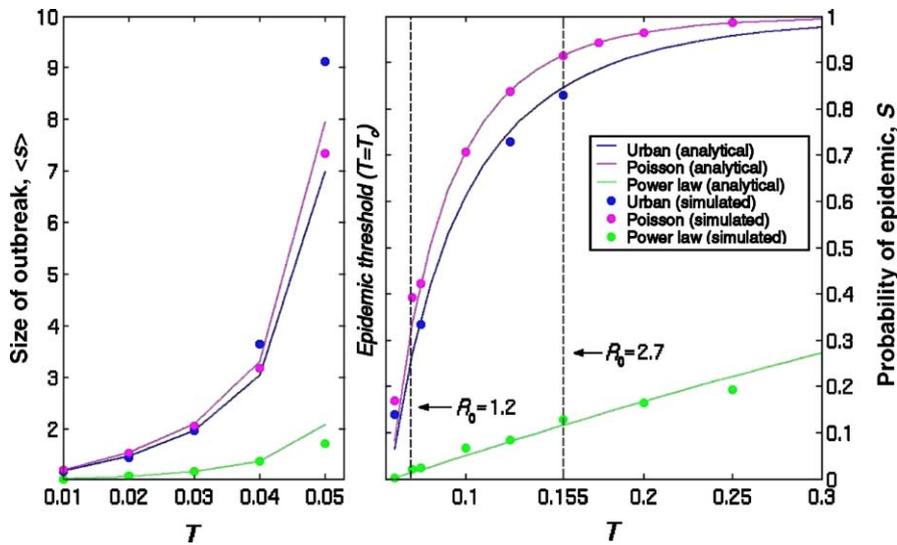


Fig. 3. Predicting outbreaks and epidemics. The left graph illustrates the average number of people infected in a small outbreak, $\langle s \rangle$, when T is below the epidemic threshold. The right graph illustrates S , the probability that an epidemic occurs when T is above the epidemic threshold. S also equals the expected fraction of the population infected during an epidemic, should one occur. The vertical lines correspond to recent estimates of R_0 for SARS pre-intervention (right) and post-intervention (left) (Chowell et al., 2003; Lipsitch et al., 2003; Riley et al., 2003). Note that we chose the parameters for the power law and Poisson networks so that for any value of T , all three networks share the same R_0 . Simulation values are based on 2571 simulated epidemics, each starting with a unique individual in the network.

estimates of R_0 for SARS in these three networks (Chowell et al., 2003; Lipsitch, 2003; Riley et al., 2003).

The different vulnerability of these networks arises from differences in the patterns of interpersonal contacts. For example, power law networks are made up primarily of vertices with very few contacts, and have a small minority of superspreaders of high degree. They will surpass the epidemic threshold when there is a non-negligible probability that an outbreak will reach members of this minority. Because superspreaders are rare, however, any particular outbreak may fail to reach this minority and thus never generate an epidemic. In contrast, vertices in a Poisson network will be fairly homogeneous and any two small outbreaks will be essentially equivalent. Whereas a power law network reaches an epidemic threshold only when transmissibility is sufficiently high to reach superspreaders with reasonable frequency, a Poisson network reaches an epidemic threshold only when the ‘typical’ outbreak leads to an epidemic. Thus in the Poisson network, most outbreaks above the epidemic threshold give rise to epidemics.

Stochastic compartmental models predict that a population with $R_0 > 1$ will experience an epidemic with probability $1 - 1/R_0$ (Bailey, 1975; Renshaw, 1991). How do we reconcile this prediction with the observation that three different networks all sharing the same R_0 have different epidemic vulnerabilities? In general, compartmental models make simple assumptions about contact patterns. The standard SIR model, for example, assumes a Poisson degree distribution. Thus, such predictions may hold for a single class of degree distribution, but cannot be generalized to arbitrary networks.

We also predict the size of a small outbreak when $T < T_c$, an estimate that is often not attainable with compartmental models (Fig. 3). Below the threshold, a typical outbreak in a power law network will die out after only one or two cases, whereas such outbreaks in the simulated urban or Poisson networks will include, on average, eight or ten individuals. Predicting the average size of a non-epidemic outbreak, $\langle s \rangle$, can aid in intervention against infectious diseases that rarely or never become self-sustaining epidemics but nevertheless have significant impact.

The similarity between the values of S and $\langle s \rangle$ predicted analytically and those measured through simulation (Fig. 3) suggest that the analytical tools provide a powerful shortcut around computationally expensive epidemic simulations. While it may not be surprising that the randomly assembled power law and Poisson networks should fit the analytical predictions well, the simulated urban network introduces additional correlations in the network structure that are not considered in our analytic calculations.

3. The impact of the initial conditions

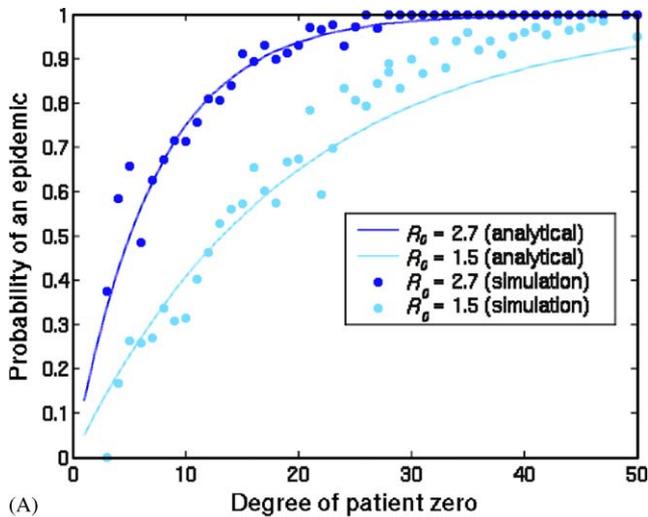
Given that, for most diseases and most communities, the introduction of a single case may or may not spark an epidemic, we can ask what factors favor one or the other outcome (epidemic or no epidemic). Using the formulas described above, we predict the likelihood of an epidemic as a function of the degree of patient zero. Both simulation and analysis of the simulated urban network illustrate that the likelihood of an outbreak is a monotonically increasing function of the degree of patient zero (Fig. 4A). Near the epidemic threshold, the risk increases exponentially with degree, whereas well above the epidemic threshold (e.g., $R_0 = 2.7$), the risk increases steeply until an epidemic is almost guaranteed. Therefore, even without differences in public health intervention, two identical communities can experience significantly different SARS outbreaks if the contact patterns of the first cases in each community differ.

In Fig. 4B, we predict the probability of an epidemic based on the size of the initial outbreak. Intuitively, the more initial cases there are, the more precarious the outcome. For diseases far above the epidemic threshold, the threat of an epidemic is overwhelming for even very small outbreaks. Thus vigilant tracking of case-contacts of the first few cases is of paramount importance in the control of such a disease.

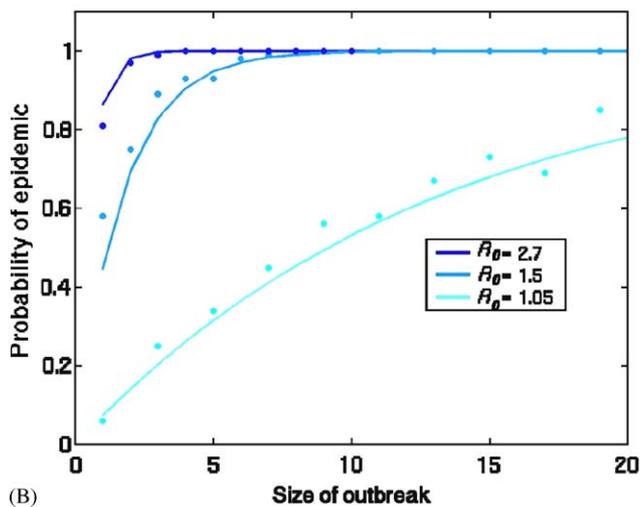
4. The impact of intervention

Along with the behavior of patient zero, individual and organized forms of disease intervention can dramatically affect an outbreak. Here we describe a simple application of our analytical toolkit to predict the efficacy of control measures. There are two basic categories of intervention. *Transmission interventions*, like wearing face masks and washing hands, reduce the likelihood that a contact with another person leads to transmission of the disease. *Contact interventions*, like avoiding public places and rearranging the patterns of interaction between caregivers and patients in a hospital, eliminate or reduce the opportunity for such contacts.

Individuals can protect themselves by adopting such measures. If followed equally rigorously, an individual will benefit equally from the two forms of intervention. That is, by reducing one’s contacts by a fraction a , an individual of degree k reduces his or her risk of infection from $1 - (1 - T + Tu)^k$ to $1 - (1 - T + Tu)^{ak}$; and reducing the likelihood of transmission per encounter by the same fraction a , reduces the risk to $1 - (1 - aT + aTu)^k$. To leading order, these two probabilities are equal. The efficacy of these strategies increases with the magnitude of the intervention a , decreases with the degree of the individual, and depends on the nature of



(A)



(B)

Fig. 4. Predicting an epidemic from initial conditions. (A) The probability that patient zero will ignite a full-blown epidemic increases monotonically with his or her degree. The calculations are based on the simulated urban network. Simulation values are based on 2571 simulated epidemics, one for every unique patient zero in the network. Discrepancy between simulations and analysis is likely caused by the finite size of the network, which contains very few high degree vertices, and by the intrinsic community structure in which high degree vertices (like teachers and caregivers) are preferentially connected to each other. (B) The probability of a full-blown epidemic in the simulated urban network increases with the size of the initial outbreak. This calculation does not assume knowledge of the specific degrees of the individuals affected in the outbreak, information that would improve the prediction. For each circle in the graph, we ran 100 simulations starting with the appropriate number of randomly selected initial cases and calculated the fraction of those outbreaks that gave rise to an epidemic.

the underlying contact network (Fig. 5). Near the epidemic threshold, individuals of different degree benefit equally from such strategies. Well above the epidemic threshold, the more limited one's contacts, the greater the impact of such interventions. For very highly

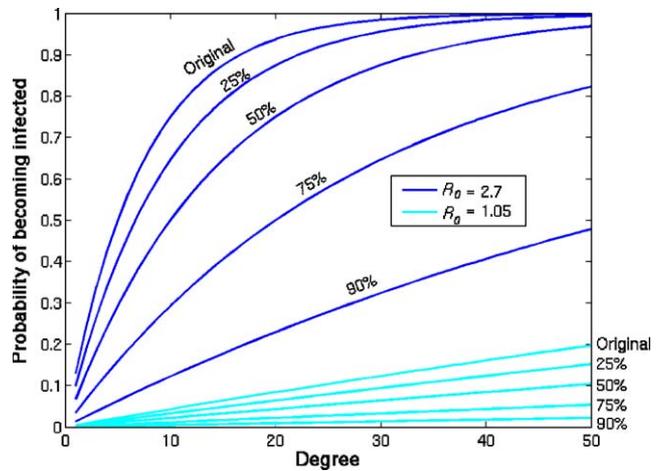


Fig. 5. Individual intervention. The probability that an individual will become infected increases with the extent of personal precautions and the average transmissibility of the disease (calculated for a simulated urban network). For example, the 25% lines indicate the risk of infection if an individual who either reduces his/her contacts by 25% or lowers the likelihood of transmission per contact by 25%.

connected individuals, partial interventions will have little or no effect.

There are situations where one form of intervention is more feasible than the other. In such cases, instead of comparing transmission and contact reductions of equal compliance (as in Fig. 5), one should compare transmission and contact interventions of equal feasibility. For example, the surgical masks seen on the streets of Hong Kong may be poorly effective against the SARS agent, and may hardly lower transmissibility, whereas the N-95 or surgical masks worn by hospital workers may be quite effective (Seto et al., 2003). Furthermore, individuals may easily reduce their social interactions by avoiding crowded shopping malls and movie theaters, whereas hospital caregivers may be unable to reduce the number of patients with whom they interact. Thus, taking into account intervention feasibility, members of the general public may be advised to avoid contact opportunities, whereas hospital employees should invest in high efficiency masks.

Policy makers take a more statistical approach to intervention than do individuals. Their goal is to reduce the likelihood and size of epidemics while minimizing cost. The most restrictive policies, such as closed borders and quarantining every member of society, are economically catastrophic. Network theory allows us to predict quantitatively the epidemiological impacts of diverse strategies, which can then be paired with economic and sociological assessments of such strategies. We provide a simple example of such analysis.

Widespread adoption of transmission interventions will lower T and may successfully bring a population

under the epidemic threshold. For example, if a disease is spreading through our simulated urban network with $T = 0.1550$ (equivalently, $R_0 = 2.7$), then Fig. 3 indicates that average transmissibility must be lowered by 62.9% to bring the population under the epidemic threshold (i.e., under $T = 0.0575$). An intervention that lowers transmission by less than 62.9% (for example, poor quality face masks) will not in itself prevent the emergence of an epidemic.

As illustrated in Fig. 6, policy makers can also consider the quantitative impact of contact reducing interventions. For a disease with $R_0 = 2.7$ ($T = 0.1662$) spreading through the Vancouver-like network, 50% reductions in contacts will have variable impact depending on the targeted demographic. Health care workers are more at risk than school children, other working adults, and non-working adults. A typical non-working adult cuts his risk of infection by approximately 33% when he or she eliminates half of previous contacts, whereas a health care worker cuts his or her risk by only 17%. These reductions also reduce the probability that the individual will spark a large-scale epidemic by similar amounts. Policy makers can consider the impact of not only individual compliance, but also of large-scale intervention. For example, closing all schools would reduce the probability of an epidemic from 84.9% to 73.8%. In a related article, we use these methods to compare a large range of control measures for a spectrum of respiratory-borne diseases (Pourbohloul et al., In review).

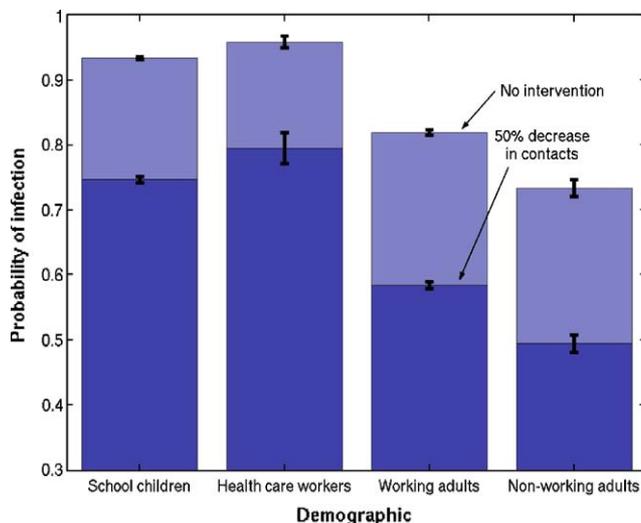


Fig. 6. Demographics of intervention. The impact of contact-reducing interventions varies by demographic sector. Health care workers are most at risk followed by school children, working adults, and non-working adults. Light bars reflect baseline risk before intervention and dark bars reflect reduced risk to an individual who limits his or her contacts to half of the previous amount. Error bars are 95% confidence intervals that reflect underlying diversity of contact patterns within each demographic sector.

5. Discussion

Contact network modeling of disease transmission allows us to make quantitative predictions about the scale of outbreaks from information about the first few cases (Fig. 4). This provides insight into the very different SARS outbreaks in Toronto and Vancouver. The first cases in both cities were exposed to SARS at Hotel M in Hong Kong on February 21, 2003. Patient zero in Toronto (onset February 25, 2003) was the matriarch of an extended, multi-generational family who died at home on March 5 as an unrecognized case of SARS (Poutanen et al., 2003). In addition to the five of ten persons within her family cluster who were affected, subsequent spread to health care workers, patients and their families culminated in more than 200 probable cases in Toronto, most of them indigenously acquired (Health Canada, 2003; Poutanen et al., 2003). In contrast, patient zero in Vancouver (onset February 26, 2003) returned from traveling with his wife to an otherwise empty abode on March 7, 2003 and was almost immediately hospitalized (Poutanen et al., 2003). Two of the three other probable cases in Vancouver were imported from abroad and no secondary case of SARS from patient zero was detected (Health Canada, 2003). Fig. 4A illustrates that, in our simulated urban network, the difference between a patient zero with one to five contacts (Vancouver) and a patient zero with ten or more contacts (Toronto) can mean a doubling or tripling in the likelihood of an epidemic.

This analysis further sheds light on the importance of contact patterns to the fate of an epidemic. Settings with different contact networks may share the same R_0 , yet may differ significantly in their vulnerability to an epidemic and the impact of an epidemic should one occur. Realistic contact networks also allow detailed quantitative assessments of intervention strategies. A caveat to this approach is the static nature of our models. While our calculations capture the temporal progression of a disease through a population, we assume that the underlying contact patterns are fixed throughout an outbreak, a reasonable assumption if public health control measures are implemented early and consistently. While one can easily consider temporally changing interventions through epidemiological simulation, we have not yet developed analytic methods to predict the impact of such dynamics.

R_0 is a valuable epidemiological quantity. It is relatively straightforward to derive based on routinely collected epidemiologic data and can be predictive in homogeneous settings; it also has historically informed effective vaccination strategies (Anderson and May, 1991; Hethcote, 2000). Yet R_0 has its limits. Since R_0 is a function of both the transmissibility of a disease and the contact patterns that underlie transmission, then measuring R_0 in a location where contact rates are unusually

high will lead to an estimate that is not appropriate for the larger community. As in the case of SARS, however, data is often only available for a limited and unrepresentative sample of the larger population.

Estimating the transmissibility T instead of R_0 gives us a way out of this difficulty. This means reporting not just the number of new infections per case, but also the total estimated number of contacts during the infectious period of that case. Given the primary role of contact tracing in infectious disease control, the relevant data is often collected. As an alternative to such direct estimates of T , if the underlying contact network is known, one can estimate T based on epidemiological case counts. We can predict the expected number of cases at each generation of transmission as a function T (not described here) and thus estimate T by curve fitting to the appropriate data (Meyers et al., 2003). The sensitivity of our epidemiological predictions to the estimated value of T will vary by both contact network and disease (Figs. 3–5). For example, for highly contagious diseases ($T > 0.2$), the probability of a large-scale epidemic is much more sensitive to minor variations in T in the power law network than in the Poisson and urban networks (Fig. 3). The reverse is true for mildly contagious diseases just above the epidemic threshold.

Unlike R_0 , T can be justifiably extrapolated from one location to another even if the contact patterns are quite disparate. We offer a simple example to illustrate the benefits of measuring T . Suppose we measure $R_0 = 2.7$ in a hospital where the average individual comes in close contact with 100 other individuals. Then the probability that an individual will catch the disease from an infected contact is just 2.7% or, in network terms, $T = 0.027$. Now suppose the typical individual in the general population has 10 close contacts that could potentially lead to the spread of a disease. If we extrapolate $R_0 = 2.7$ to the general public, then we predict that, on average, 2.7 out of every 10 contacts or 27% of contacts become infected, whereas if we extrapolate $T = 0.027$ to the general public we still have only 2.7% of contacts becoming infected, which gives us a much reduced expectation for the spread of the disease.

We have illustrated that percolation theory allows us to shortcut time-consuming simulations to produce robust predictions about the size and likelihood of an epidemic, the implications of patient zero's contact patterns and the initial size of an outbreak, and optimally reducing the risk of infection at a personal and population level. To truly reap these benefits of the network epidemiology, we must have not only good estimates of T but also realistic models of contact networks. We have presented a first step towards a realistic community network and have shown that it departs significantly from the idealized networks previously used in network epidemiology, and yet still is

sufficiently random for the application of epidemiological network analysis. The task of making even more realistic contact network models of small communities, hospitals, cities, even the global the population is enormous, but promises a step change in our ability to predict and effectively control the spread of infectious disease. As we incorporate better data at these various scales, network theory will allow us to generalize our predictions and make better suggestions for epidemiological control.

Acknowledgements

The authors thank J.J. Bull, M. Girvan, and S. Meyers for useful conversations and editorial suggestions. This work was supported in part by grants from Canadian Institutes of Health Research (CIHR) to the BC SARS Investigators Collaborative that includes L.A.M., B.P., D.M.S., and R.C.B., from the National Science Foundation (DEB-0303636) to L.A.M., and from the James S. McDonnell Foundation and the National Science Foundation (DMS-0234188) to M.E.J.N. The Santa Fe Institute and CIHR supported the working visits of B.P., L.A.M., and M.E.J.N.

References

- Anderson, R.M., May, R.M., 1991. *Infectious Diseases of Humans, Dynamics and Control*. Oxford University Press, Oxford.
- Bailey, N.T.J., 1975. *The Mathematical Theory of Infectious Diseases and Its Applications*. Hafner Press, New York.
- Ball, F., Mollison, D., Scalia-Tomba, G., 1997. Epidemics with two levels of mixing. *Ann. Appl. Probab.* 7, 46.
- Barabasi, A.L., Albert, R., 1999. Emergence of scaling in random networks. *Science* 286, 509.
- BC Stats, 2002. Labour Force—Employment, Unemployment, Related Rates, BC Stats.
- BC Stats, 2003. 2001 Census Profile of British Columbia's Regions: Greater Vancouver Regional District, BC Stats.
- Becker, N., 1977. Estimation for discrete time branching processes with application to epidemics. *Biometrics* 33, 515–522.
- Booth, C. M., Matukas, L. M., et al., 2003. Clinical Features and Short-term Outcomes of 144 Patients With SARS in the Greater Toronto Area. *JAMA*, 289.21.JOC30885.
- Centers for Disease Control and Prevention, 2003. Severe Acute Respiratory Syndrome (SARS) Report of Cases in the United States, Centers for Disease Control and Prevention.
- Centre for Health Services and Policy Research, 2002. *The British Columbia Health Atlas*, Centre for Health Services and Policy Research.
- Chowell, G., Castillo-Chavez, C., et al., 2004. Model Parameters and Outbreak Control for SARS. *Emerging Infectious Disease* 10.
- Chowell, G., Fenimore, P.W., et al., 2003a. SARS outbreaks in Ontario, Hong Kong and Singapore: the role of diagnosis and isolation as a control mechanism. *J. Theor. Biol.* 224, 1–8.
- Chowell, G., Hyman, J.M., et al., 2003b. Scaling laws for the movement of people between locations in a large city. *Phys. Rev. E* 68, 066102.
- Cyranoski, D., Abbott, A., 2003. Virus detectives seek source of SARS in China's wild animals. *Nature* 423, 467.

- Diekmann, O., De Jong, M.C.M., Metz, J.A.J., 1998. A deterministic epidemic model taking account of repeated contacts between the same individuals. *J. Appl. Probab.* 35, 448–462.
- Donnelly, C. A., Ghani, A. C., et al., 2003. Epidemiological determinants of spread of causal agent of severe acute respiratory syndrome in Hong Kong. *Lancet*, 1.
- Drosten, C., Gunther, S., et al., 2003. Identification of a Novel Coronavirus in Patients with Severe Acute Respiratory Syndrome. *N. Engl. J. Med.* 348, 1867.
- Durrett, R., 1999. Stochastic spatial models. *SIAM Rev.* 41, 577–718.
- Eubank, S., Guclu, H., et al., 2004. Modelling disease outbreaks in realistic urban social networks. *Nature* 429, 180–184.
- Farrington, C., Kanaan, M., Gay, N., 2003. Branching process models for surveillance of infectious diseases controlled by mass vaccination. *Biostatistics* 4, 279–295.
- Ferguson, N.M., Garnett, G.P., 2000. More realistic models of sexually transmitted disease transmission dynamics: sexual partnership networks, pair models, and moment closure. *Sex. Transm. Dis.* 27, 600.
- Greenhalgh, D., 1986. Optimal control of an epidemic by ring vaccination. *Comm. Stat.: Stoch. Models* 2, 339–363.
- Health Canada, 2003. Summary of Severe Acute Respiratory Syndrome (SARS) Cases: Canada and International, April 28, 2003. http://www.hc-sc.gc.ca/pphb-dgspsp/sars-sras/eu-ae/sars20030428_e.html.
- Hethcote, H., 2000. Mathematics of infectious diseases. *SIAM Rev.* 42, 599.
- Hethcote, H.W., Yorke, J.A., 1984. *Gonorrhea Transmission Dynamics and Control*. Springer, Berlin.
- Keeling, M.J., Rand, D.A., Morris, A.J., 1997. Correlation Models for Childhood Diseases. *Proceedings of the Royal Society (London) B* 264, 1149–1156.
- Keeling, M.J., Woolhouse, M.E., et al., 2003. Modelling vaccination strategies against foot-and-mouth disease. *Nature* 421, 136–142.
- Kleczkowski, A., Grenfell, B.T., 1999. Mean field-type equations for spread of epidemics: the ‘small world’ model. *Physica A* 274, 1–385.
- Ksiazek, T.G., Erdman, D., et al., 2003. A Novel Coronavirus Associated with Severe Acute Respiratory Syndrome. *N. Engl. J. Med.* 348, 1953.
- Lefevre, C., Picard, P., 1989. On the formulation of discrete-time epidemic models. *Math. Biosci.* 95, 27–35.
- Leo, Y.S., Chen, M., et al., 2003. Severe acute respiratory syndrome—Singapore, 2003. *Morb. Mort. Wkly. Rep.* 52, 405.
- Lipsitch, M., Cohen, T., et al., 2003. Transmission Dynamics and Control of Severe Acute Respiratory Syndrome. *Science*, 1086616.
- Lloyd, A.L., May, R.M., 2001. Epidemiology. How viruses spread among computers and people. *Science* 292, 1316.
- Longini, I.M., 1988. A mathematical model for predicting the geographic spread of new infectious agents. *Math. Biosci.* 90, 367.
- Marra, M.A., Jones, S.J.M., et al., 2003. The Genome Sequence of the SARS-Associated Coronavirus. *Science* 300, 1399.
- Maslov, S., Sneppen, K., Zaliznyak, A., 2001. Pattern detection in complex networks: Correlation profile of the internet. Preprint cond-mat/025379.
- Meyers, L.A., Newman, M.E.J., et al., 2003. Applying network theory to epidemics: control measures for *Mycoplasma pneumoniae* outbreaks. *Emerg. Infect. Dis.* 9, 204.
- Morris, M., 1995. Data driven network models for the spread of disease. In: Mollison, D. (Ed.), *Epidemic Models: their Structure and Relation to Data*. Cambridge, Cambridge University Press, pp. 302–322.
- Müller, J., Schönfisch, B., Kirkilionis, M., 2000. Ring vaccination. *J. Math. Biol.* 41, 143–171.
- Newman, M.E.J., 2002. Spread of epidemic disease on networks. *Phys. Rev. E* 66, art. no.-016128.
- Pastor-Satorras, R., Vespignani, A., 2001. Epidemic spreading in scale-free networks. *Phys. Rev. Lett.* 86, 3200.
- Peiris, J.S.M., Lai, S.T., et al., 2003. Coronavirus as a possible cause of severe acute respiratory syndrome. *Lancet* 361, 1319.
- Pourbohloul, B., Meyers, L. A., et al., In review. Halting the spread of respiratory infections: A quantitative comparison of prevention and control strategies.
- Poutanen, S.M., Low, D.E., et al., 2003. Identification of Severe Acute Respiratory Syndrome in Canada. *N. Engl. J. Med.* 348, 1995.
- Renshaw, E., 1991. *Modelling Biological Populations in Space and Time*. Cambridge University Press, Cambridge, UK.
- Riley, S., Fraser, C., et al., 2003. Transmission Dynamics of the Etiological Agent of SARS in Hong Kong: Impact of Public Health Interventions. *Science*, 1086478.
- Ritton, T., O’neill, P.D., 2002. Bayesian Inference for Stochastic Epidemics in Populations with Random Social Structure. *Scand. J. Stat.* 29, 375–390.
- Sander, L.M., Warren, C.P., et al., 2002. Percolation on heterogeneous networks as a model for epidemics. *Math. Biosci.* 180, 293–305.
- Sattenspiel, L., Simon, C.P., 1988. The spread and persistence of infectious diseases in structured populations. *Math. Biosci.* 90, 341.
- Seto, W.H., Tsang, D., et al., 2003. Effectiveness of precautions against droplets and contact in prevention of nosocomial transmission of severe acute respiratory syndrome (SARS). *Lancet* 361, 1519.
- Statistics Canada, 2001. Household size, census metropolitan areas, Statistics Canada.
- Van Der Ploeg, C.P.B., Van Vliet, C., et al., 1998. A microsimulation model for decision support in STD control. *Interfaces* 28, 84–100.
- Vancouver School Board, 2002. Vancouver School Board December 2002 Ready Reference, Vancouver School Board.
- World Health Organization, 2003. Severe Acute Respiratory Syndrome (SARS), World Health Organization.
- Xu, R.-H., He, J.-F., et al., 2004. Epidemiologic Clues to SARS Origin in China. *Emerg. Infect. Dis.* 10.
- Yu, I.T.S., Li, Y., et al., 2004. Evidence of airborne transmission of the severe acute respiratory syndrome virus. *N. Engl. J. Med.* 350, 1731–1739.