Infection Dynamics on Small-World Networks

Alun L. Lloyd, Steve Valeika, and Ariel Cintrón-Arias

Abstract. The use of network models to describe the impact of local spatial structure on the spread of infections is discussed. In particular, we focus on small-world networks, within which the pattern of interactions can be varied from being entirely local to being entirely global as a single parameter is changed. Analysis approaches from graph theory, statistical physics and mathematical epidemiology are discussed. Simulation results are presented that highlight the surprising findings of Watts and Strogatz [**59**], namely that a small number of long-range interactions in an otherwise locally structured population can markedly enhance the ability of an infection to spread and the rate at which the spread occurs. We also discuss some of the implications of such spatial structure on the dynamics and persistence of endemic infections.

1. Introduction

It has long been realized that the spatial structure of a population can have a major impact on the spread of infectious diseases. While simple, non-spatial, mathematical models have given many insights into the dynamics of transmission, many situations call for more realistic models that include some description of space.

The simplest epidemic models assume that the population is well-mixed, so that any pair of individuals is equally likely to interact with each other during a given time interval. Perhaps the simplest way of extending the model to account for spatial structure is to subdivide the population into two or more 'patches' [**10, 41**], representing, for instance, different cities. It is typically assumed that these patches are well-mixed and that there is some level of mixing between different patches. The between-patch mixing usually occurs at a much lower rate than within-patch mixing. Such models are often termed metapopulation models.

An alternative way to describe spatial structure assumes that individuals are distributed continuously across space and that infection spreads as individuals move about the landscape. A random-walk description is often employed for this movement, leading to the appearance of diffusion terms in the model: a so-called reaction-diffusion model [**50**]. Another spatially continuous formulation assumes

1991 *Mathematics Subject Classification.* Primary 92D30; Secondary 60K35.

A.L. and S.V. were supported by the University of North Carolina Center for AIDS Research (CFAR), an NIH-funded program (P30 AI50410).

A. C.-A. was supported by the Department of Mathematics and Statistics of Arizona State University, and the Statistical and Applied Mathematical Sciences Institute (SAMSI), Research Triangle Park, NC, an NSF-funded center (DMS-011209).

that individuals are fixed in space but are able to infect individuals at other locations, with the probability of transmission depending on the geographic separation being described by an infectivity kernel [**37**]. Study of these model types has focused on the speed at which infection invades a population as it spreads outwards from its initial point of introduction in a wave-like fashion [**34, 57**].

Metapopulation and spatially continuous models have dominated much of the spatial epidemiology literature. These population-level models are typically cast as either a set of coupled ordinary differential equations, in the case of metapopulation models, or a partial differential equation or integro-differential equation, in the case of spatially continuous models. Such models can be amenable to mathematical analysis.

Network models [**46**] provide a quite different, individual-based, approach to studying the impact of spatial structure. The members of the population are modeled as the nodes of a network, and the edges of the network represent interactions between people that could potentially lead to transmission of the infection. The complexity of network models arises because they must account for each individual in the population as well as describing how all of these individuals interact with each other. The need to fully describe the interaction network is a major difficulty for the practical application of network approaches.

An important use of network models has been as a research tool, providing a framework within which the impact of the detailed assumptions of spatial structure that underlie many of the population-level models can be explored. An issue of particular interest is the way in which interactions that are specified in terms of the behavior of individuals impact upon population-level behavior [**17, 32**], often discussed under the banner of 'from individuals to populations'. With this approach, one hopes to identify which features of individual-based models have an important effect at the population level.

Much progress has been made by employing simple classes of networks that capture particular aspects of population structure. For instance, in the case of spatial structure, one might contrast settings in which mixing is mainly local in nature against settings in which the population is essentially well-mixed. One might then ask what happens in intermediate cases, where most of the mixing is local but there are occasional long-range interactions.

While networks that exhibit purely local or purely global mixing have long been studied, the introduction of small-world networks in a seminal paper by Watts and Strogatz [**59**] enables the exploration of intermediate settings. The interest created by this paper has led to a flourishing literature on network approaches in a diverse range of settings.

This chapter reviews the recent literature on small-world networks as it relates to epidemiological modeling. In Section 2 we introduce small-world networks, and review their properties. Section 3 discusses the impact of small-world networks on basic properties of epidemics, such as disease transmission thresholds. Section 4 considers the dynamics of epidemics, including the timecourse of outbreaks, the probability of their occurrence and the size of the resulting outbreak. Section 5 considers dynamical issues in endemic settings, including the impact of stochasticity on behavior near the equilibrium and the possibility of oscillatory behavior.

2. Construction and Properties of Small-World Networks

Regular lattices and random graphs both have a long history of use in network theory and as models for the structure of populations. A classic example of a lattice model is provided by Harris's contact process model [**24**]. Lattice models assume that individuals are located at the sites of a regular lattice and connections are made to some collection of the nearest neighbors of each site. As an example, in one dimension, individuals may be regularly spaced along a line and each is assumed to interact with their k nearest neighbors. In two dimensions, individuals might be sited on a regular square grid, with connections made to their four nearest neighbors (up, down, left and right: the so-called von Neumann neighborhood) or their eight nearest neighborhood (up, down, left, right, and the four diagonals: the Moore neighborhood). In order to avoid edge effects, periodic boundary conditions are often imposed. In our one dimensional example, the line would be wrapped around onto a circle so that the first and last individuals on the line would become neighbors.

The random graph [**11**], most commonly associated with Erdös and Renyi from their graph theoretical treatment of this setting, assumes that each pair of individuals has some probability, q, of being connected. These connections are made randomly and independently between the pairs. The major difference between regular lattices and random graphs is that interactions are purely local in the former—individuals are only connected to their neighbors—whereas interactions are purely global in the latter: connections are made with no regard for the spatial location of individuals.

Small-world networks, introduced by Watts and Strogatz [**59**], are intermediate between the regular lattice and a random graph. Watts and Strogatz produced these networks by starting from a regular lattice and randomly rewiring a certain proportion, p, of the network's links. In this paper, we employ a rewiring process in which both ends of the link are moved and connected randomly to other individuals in the population[1]. These rewired edges are termed long-range connections: they are made without any regard to spatial location and so will typically not be local in nature. When the rewiring probability is zero, the process leaves the lattice unaltered. When the rewiring probability approaches one, all of the links are rewired and so have no dependence on spatial location. The resulting network is, in many ways, similar to the random graph.

We remark that a special case of the Watts and Strogatz small-world network was introduced in an earlier paper by Ball et al. [**4**]. Their 'great circle' model locates individuals on a circle, with local connections being made to the closest neighbors on the left and right, each with probability q_{L}, and global connections being made between any pair in the population with probability q_{R}. Ball et al. make the remark that their model could be generalized.

An attractive feature of the original Watts and Strogatz algorithm for generating small-world networks (and of the slightly altered algorithm that we employ here) is that it preserves the total number of connections present in the original

[1]In their paper, Watts and Strogatz only moved one end of each rewired link. A systematic method was used to determine which end of the link was moved: for their one dimensional circular lattice, they looked at links in a clockwise sense and kept the first end of the link fixed. Each node, therefore, would retain at least half of its original links, even if all links in the network are rewired.

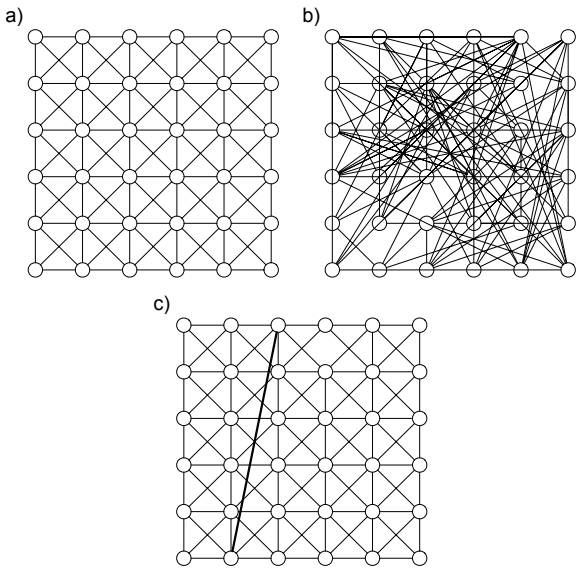

FIGURE 1. Example networks. (a) Two dimensional lattice with connections to the eight nearest neighbors (Moore neighborhood). For clarity, we do not impose periodic boundary conditions. (b) Random graph, with, on average, connections made to 8 neighbors of each site. (It should be noted that not all of the connections are visible: many of links are shown superimposed.) (c) Small-world network obtained by randomly rewiring a fraction of the links of the lattice. In this example, for which the rewiring probability was low, just a single link has been rewired. Rewiring typically leads to the appearance of long-range links.

lattice. This allows for a fair comparison to be made between the lattice and the generated graph. Another variant of the generating algorithm involves the addition of extra links, rather than the rewiring of existing links. It turns out that this latter algorithm is mathematically a more satisfactory way of generating small-world networks. The fact that additional links are being added to the network must be kept in mind: mixing is inherently greater in this latter small-world network, although this effect is negligible if the probability of link addition is low.

2.1. Network Properties. In order to make precise the way in which small-world networks fall between regular lattices and random graphs, we must first introduce some of the measures by which networks are described.

A **connected** network is one for which infection could, in theory, travel from any person in the population to any other person in the population. Such transmission will typically involve a chain of infections, involving a number of intermediate individuals. The lengths of these transmission chains are captured by various notions of distances in the network. The **distance** between two individuals is defined as the length of the shortest path by which one can move from one individual to the other. These distances can be summarized by a quantity such as the **average distance**, taken over all pairs of individuals, which gives an idea of the typical

distance between individuals [**14, 59**], or the **diameter** of the network, which is the largest of the distances, taken over all pairs of individuals.

The number of (immediate) neighbors of a given individual is known as their **degree** or **connectivity** and is typically denoted by k. Looking across the entire population, these quantities give rise to the **connectivity distribution** (or degree distribution). This distribution is often described in terms of its mean, often written as $\langle k \rangle$, and variance. Positive values of this variance correspond to heterogeneous degree distributions. Much of the literature on network dynamics has focused on the impact of hetereogeneity in the degree distribution [**33, 43, 46, 51**]. Many real-world networks exhibit marked variation in the connectivities of different individuals, with, in many cases, the variance being much larger than the mean. Such networks are often simply called heterogeneous networks, but since degree hetereogeneity is not the only form of heterogeneity that a network can exhibit, this abbreviated term should be used with some care if the context is not totally clear.

Various measures of **cliqueishness** [**26, 44, 45, 49, 59**], including the clustering coefficient, transitivity and mutuality, examine the extent to which the neighborhoods of connected individuals overlap. Cliques of size three arise when two connected individuals have a common neighbor, leading to the appearance of triangles in the network. The presence of such cliques is captured by the **transitivity**, ϕ, of the network.[2] All of the triples in the network (i.e. paths of length three: instances in which individual A is connected to individual B who is connected to individual C) are examined and ϕ is calculated as the fraction of these that close up into triangles (i.e. those for which individual A is also directly connected to individual C) [**26, 45**].

Clearly ϕ does not capture all aspects of cliqueishness, as illustrated by the case of a two dimensional lattice with the von Neumann neighborhood. This network contains no triangles, so ϕ is equal to zero. Instead, its cliques involve loops of length four, a property captured by the **mutuality** [**44**] of the network. More generally, the notion of cliques can be extended to cliques of various sizes, corresponding to the existence of short length loops of varying lengths.

Regular lattices are connected networks. A consequence of the local nature of the interactions in a regular lattice is that path lengths are long. For regular d-dimensional cubic lattices with linear dimension L, so that $N = L^d$, path lengths scale linearly with L, or, in terms of the population size, scale as $N^{1/d}$ [**48, 59**]. Typical path lengths are shorter when individuals are more highly connected, with path lengths scaling with the reciprocal of the linear dimensions of the neighborhood [**48, 59**]. This means that path lengths scale with the reciprocal of the the connectivity if the neighborhood consists of the k nearest neighbors on a one dimensional lattice or in a k-sized von Neumann neighborhood in two dimensions (i.e. individuals are connected to their closest $k/4$ neighbors in the up, down, left and right directions). The local structure also means that regular lattices exhibit high levels of cliqueishness (see figure 2c). (As an example, for the lattice that we employ in our epidemic simulations—see the caption of figure 3 for details of its

[2]Some authors refer to ϕ as the clustering coefficient. Unfortunately, there is another measure of cliqueishness that is also known as the clustering coefficient [**59**]. Consequently, to minimize confusion, we avoid using the term here.

construction— it is straightforward to show that $\phi = 2/3$.) All individuals in a lattice (excepting those on the edge of the lattice if periodic boundary conditions are not imposed) have the same number of neighbors: the connectivity distribution is homogeneous.

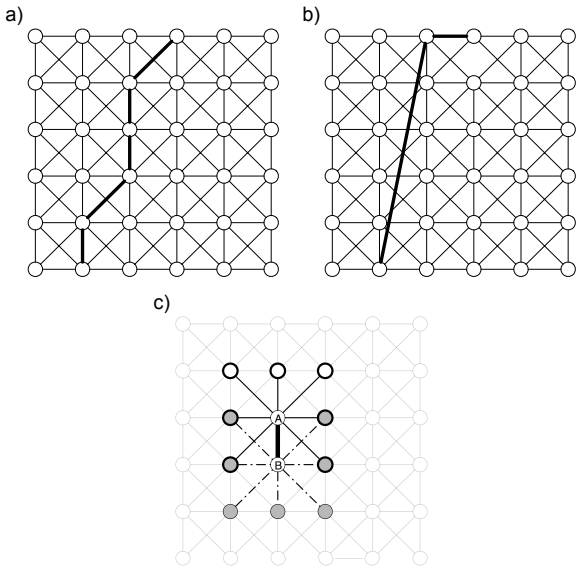

FIGURE 2. Network Properties. (a) and (b) Path lengths are long in the regular lattice, but are rapidly shortened with the inclusion of long-range links. (c) Cliques are common in regular lattices. In this example, we focus on two individuals, A and B, who are connected. We see that four individuals are neighbors of both A and B. Notice that these cliques correspond to the occurrence of triangles in the network.

Random graphs are not necessarily connected. If the connection probability, q, is small then the graph is typically composed of a large number of small, disconnected components [**11**]. The distribution of the sizes of these components is exponential with finite mean, even as N tends to infinity. The typical size of these components is $O(1)$: their average size does not scale with N (provided that N is sufficiently large)[3]. When q is sufficiently large, the random graph typically consists of one connected subgraph that includes a large fraction of the population, together with a number of small disconnected components. This component is known as the 'giant component' of the graph and its size is $O(N)$ (that is to say, the size of the giant component scales linearly with N: notice that its average number of nodes, therefore, diverges as $N \to \infty$, but that average fraction of nodes in the giant component approaches a constant). The size distribution of the remaining small components is exponential, with finite mean. Again, the typical size of these is $O(1)$. A well-known theorem [**11**] makes these statements more precise, stating

[3]The size of the **largest** of the small components scales with the logarithm of N. This is because the number of small components increases as the number of nodes increases, so the tail of the size distribution is explored more fully.

that (for $N \to \infty$) the random graph has a (single) giant component if and only if $\Phi = Nq$ is greater than one. (We shall see that Φ is simply the average connectivity of the nodes in the network). This component then contains a proportion z of the population, where z is the greatest root of the equation

$$z = 1 - \exp(-\Phi z). \tag{1}$$

The global nature of mixing in random graphs means that distances are short compared to those in regular lattices. Path lengths scale with the logarithm of the population size and with the reciprocal of the logarithm of the average connectivity [**59**]. Cliques are rare in random graphs, with ϕ scaling as $1/N$ for large N [**46, 59**]. The connectivities of the nodes of the random graph are binomially distributed, according to $B(N-1, q)$. For large N this distribution approaches a Poisson distribution with mean $(N-1)q \approx Nq$. Since the variance and mean are equal for Poisson distributions, the variance of the connectivity distribution equals Nq. We remark that the connectivity distribution of the random graph is not very heterogeneous: few nodes have connectivity that differ that greatly from the average.

For small values of the rewiring parameter, p, (or for a small frequency of additional links) most of the links in the small-world network are simply those of the lattice and there are only a few long range connections. The surprising result of Watts and Strogatz is that these few long-range connections rapidly shorten path lengths in the network. As the parameter p is increased, path lengths quickly fall to become comparable to those in the random graph [**59**]. Cliqueishness, on the other hand, is much less affected by rewiring: it is not until p takes values approaching one that the remnants of the highly clustered lattice are destroyed by rewiring. For a wide range of p values, the Watts and Strogatz algorithm generates networks that have the short path length of the random graph while having the high level of cliqueishness of the lattice: this is the small-world regime. Small-world networks, in this sense, exhibit both local and global mixing properties.

The heterogeneity of the small-world networks generated using the Watts-Strogatz method is intermediate between those of the lattice and the random graph. In the $p = 1$ case, the degree distribution of the Watts-Strogatz graph generated using their original algorithm (in which only one end of each link is rewired, with the fixed end being chosen systematically) has a lower variance than that of the Poisson distribution of the random graph [**6**]. In this respect, their totally rewired graph differs from the random graph[4]. The rewiring algorithm that we employ, in which both ends of links are rewired, leads to a degree distribution that is closer to that of the random graph. (The resulting graph is still not quite the Erdös-Renyi random graph described earlier, since the total number of links in the network is fixed.) In none of these cases, however, is the variance particularly large and so heterogeneity of the degree distribution will not play a major role in the dynamics of infections on the graphs considered here.

We remark that many networks with heterogeneous degree distributions also exhibit short average path lengths [**14**]. Here, we restrict use of the term 'small-world network', and our attention, to the Watts-Strogatz type of small-world network.

[4]Also recall the earlier comment that the original Watts and Strogatz algorithm guarantees that each node has connectivity at least $\langle k \rangle/2$. The resulting network—even when fully rewired—is guaranteed to remain connected. In contrast, the rewiring algorithm that we employ in this paper can lead to nodes becoming disconnected.

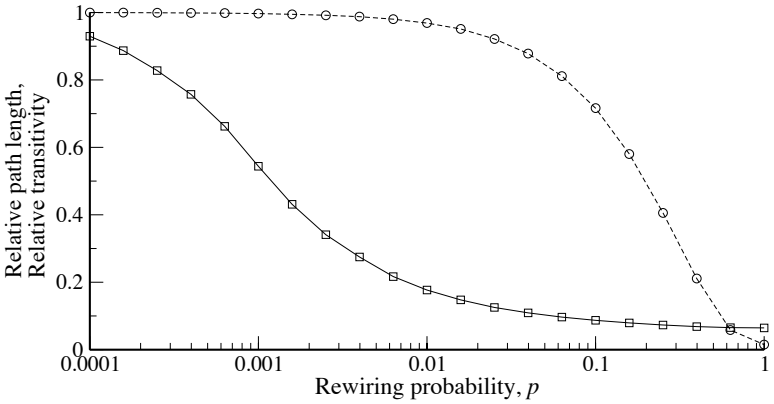

FIGURE 3. Path lengths and cliqueishness in small-world networks. Small-world networks were generated using the Watts and Stogatz method [**59**]. Starting from a one dimensional lattice of 1000 individuals, each of which were connected to their ten nearest neighbors, links were rewired using the rewiring technique discussed in the text. Periodic boundary conditions were assumed. The solid curve with squares shows the dependence of the average path length, relative to the average path length of the unrewired lattice, on the per-link rewiring probability, p. The broken curve with circles shows the cliqueishness, as measured by the transitivity (ϕ), relative to that of the unrewired lattice, which equals 2/3. Each point on the figure represents the average value taken over 200 realizations of the rewired network.

3. Analysis Techniques and Basic Properies

We first consider a particularly simple infection process. Individuals are assumed to be initially susceptible to the infection. An infectious individual can transmit infection to a susceptible; we assume that, once infected, the person is infectious immediately. Over time, infectious individuals can recover and it is assumed that recovery confers permanent immunity to the infection. This description of infection is known as the SIR (susceptible/infectious/recovered) process [**1, 15**].

It is typically assumed that there is a constant rate of transmission between an infective and susceptible who are in contact, and that this rate is the same for all such infective/susceptible pairs. Writing this transmission rate as β, we have that the probability of transmission in the short time interval $(t, t + dt)$ is equal to $\beta\, dt$. The recovery process can be described in many different ways, but most often it is either assumed that recovery occurs at a constant rate, γ, or that recovery occurs at some fixed time, τ, after infection. The constant recovery rate assumption leads to the distributions of infectious periods being exponentially distributed, with mean equal to $1/\gamma$. In order to compare the two descriptions of recovery, the values of τ and $1/\gamma$ are taken to be the same.

3.1. Well-Mixed Population-Level Models. A familiar population-level description of the SIR process in a well-mixed population, assuming a constant rate

of recovery, is given in terms of the following set of differential equations [**1, 15**]

$$\dot{x} = -cbxy \tag{2}$$

$$\dot{y} = cbxy - \gamma y \tag{3}$$

$$\dot{z} = \gamma y. \tag{4}$$

Here x, y and z denote the fractions of the population who are susceptible, infectious and recovered, respectively. In this formulation, the parameter c is the rate at which an individual makes contacts with others and the parameter b is the probability that a given contact (if made between an infective and susceptible) would lead to transmission of the infection. This formulation assumes that the population is closed: no individuals leave or enter the population, so one need not worry about demographic processes. As a consequence, equation (4) is redundant: because an individual is either susceptible, infectious or recovered, one can calculate z as $1 - x - y$.

The behavior of this deterministic model is well known [**1, 15**]. Since there is no replacement of susceptibles, introduction of infection either leads to a single epidemic, which is self-limited due to the ensuing depletion of susceptibles, or no epidemic occurs. A threshold condition governs the occurrence of these alternatives: if the value of the so-called basic reproductive number, R_0, is greater than one then an epidemic can occur, otherwise the number of infectives can never increase.

The basic reproductive number can be written in terms of model parameters as

$$R_0 = cb/\gamma, \tag{5}$$

and has a simple epidemiological interpretation. For the initial stages of the outbreak, when almost the entire population is susceptible, the parameter c gives the rate at which an individual, in particular an infective person, encounters susceptible individuals. Since $1/\gamma$ is the average duration of infection, c/γ gives the average number of susceptibles encountered over their infectious period. As the product of the average number of susceptibles encountered and the transmission probability, R_0 gives the average number of secondary infections that arise when an infectious individual is introduced into an otherwise entirely susceptible population. In this well-mixed deterministic setting, and in the absence of demography, the value of the basic reproductive number does not depend on the distribution of infectious periods: R_0 reflects the average number of secondary infections rather than the timing of their occurrence.

The total fraction of the population who become infected over the course of the entire epidemic, which we call the size of the epidemic and write as y_∞, is given by the largest root of the following equation

$$y_\infty = 1 - \exp\left(-R_0 y_\infty\right). \tag{6}$$

If R_0 is greater than one then this quantity is positive.

In the preceding discussion, we deliberately used different parameterizations for the transmission processes in well-mixed settings and in network settings. Many studies have compared infection dynamics in these two settings, in which case one must be able to move between the two parameterizations in order to make comparisons. In such studies, the parameter combination cb is typically identified with βk [**26, 28**]. Use of the well-mixed description, together with this identification of parameters, leads to the expression $R_0 = \beta k/\gamma$, or $R_0 = \beta k\tau$, for the basic

reproductive number for random networks[5]. As we shall see below, more careful consideration shows that this expression is incorrect in a couple of important ways for network settings.

3.1.1. *Stochastic Models in Well-Mixed Settings.* The deterministic models of the previous section treated the numbers of susceptible, infectious and recovered individuals as continuously varying quantities, whose changes could be described by a set of differential equations. In reality, the numbers of individuals of different types are integers, and change discretely as infection or recovery events occur. The finite size of a population and the ensuing stochastic effects (known as demographic stochasticity) can be accounted for using stochastic formulations of the well-mixed model (see, for example, [**3**]).

The stochastic well-mixed model also exhibits threshold effects as the basic reproductive number, which is again given by equation (5), is increased from below to above one [**3, 15**]. Below the threshold, each infective gives rise to an average of fewer than one secondary infection and so introduction of a single infective (or a small number of infectives) can only give rise to a minor outbreak. It should be noted that some of these 'minor' outbreaks can involve a significant number of individuals: even though the average number of secondary infections is less than one, chance events can lead to a few individuals having many more secondary infections than this.

Above the threshold, introduction of infection can lead to a major outbreak, potentially affecting a large fraction of the population. Large outbreaks are not, however, guaranteed to occur since it is possible for an infective to recover before passing on the infection. Stochastic extinction can occur if this happens for all of the infectives present at some point in time: a minor outbreak will occur if this happens early after the introduction of infection.

Using branching process theory, expressions have been derived for the probabilities of the occurrence of major and minor outbreaks when R_0 is greater than one (see, for example, [**15**]). The most familiar result states that, for the constant recovery rate model, the probability of a major outbreak occurring following the introduction of a single infective is $1 - (1/R_0)$. If recovery is instead assumed to occur at exactly time τ after infection, the probability of a major outbreak, π, is given by the largest root of

$$\pi = 1 - \exp(-R_0\pi). \tag{7}$$

3.2. Percolation and Graph Theory Approaches. The similarity between equation (1) for the size of the giant connected component of a random graph and equation (6) for the size of an epidemic in a well-mixed population (and, indeed, equation (7) for the probability of a major outbreak) is no coincidence. The corresponding threshold conditions for the connectedness of the random graph (average connectivity greater than one) and the $R_0 = 1$ threshold (average number

[5]This expression ignores all sources of heterogeneity in the network, including heterogeneity in the connectivity distribution. Degree heterogeneity can be accounted for in the differential equation framework by subdividing the population into subgroups according to individuals' connectivities. Use of this approach has led to the development of analogous expressions for R_0 that account for heterogeneity in the degree distribution [**33, 51**].

of secondary infections greater than one) are also similar to each other[6]. Noted by Barbour and Mollison [**5**], questions concerning epidemic processes on networks can be rephrased in terms of questions about the properties of graphs. As Barbour and Mollison point out, this means that the body of theory developed to describe the properties of graphs is informative about epidemic processes that occur on those graphs. One such approach that has been fruitfully applied is percolation theory.

The bond percolation problem on a graph [**23**] assumes that each edge in the graph can independently be traversed with some probability, q. Percolation theory then addresses questions such as the extent to which the network can be traversed. The earliest percolation studies focused on regular lattices and so the nature of percolation on such lattices has been described quite fully [**23**]. Following the introduction of small-world networks, their percolation properties have been characterized in detail [**38, 39, 42, 47**].

We note that the construction of a random graph can be described in terms of the bond percolation problem by starting with a completely connected graph and then taking the traversal probability of each edge of the connected graph to be q.

Grassberger [**22**], assuming that transmission rates were equal along each edge in the network and that individuals had a fixed duration of infection, noted that the spread of infection on a graph could also be mapped onto the bond percolation problem. In this simple epidemiological setting the correspondence with bond percolation is intuitively clear: transmission along an edge occurs with some probability, known as the transmissibility, which we then interpret as the traversal probability in the bond percolation model. We can then follow the infection spreading across the network and ask about the size of the network component that can be reached from the initial infective if one follows edges along which transmission occurs. (We remark that not all traversable edges in the bond percolation model will correspond to transmission events. Only traversable edges that can be reached from the initial infective can be interpreted in this way; the extra edges are not of interest in the infection context.)

Sander and co-workers [**55**] generalized this approach to a broad class of epidemic models. Newman [**43**] gives a lucid description of the correspondence between epidemic spread and percolation. Since the percolation problem is, for a given network, formulated solely in terms of the probability of transmission along edges (assuming that one node is infective and the other susceptible), the rates of transmission along edges and the durations of infection of nodes need not all be the same. Newman argues that percolation approaches can be applied to any setting in which these quantities are independent, identically distributed quantities, since then the probabilities of transmission are, *a priori*, equal along all edges. In such situations, the transmissibility is calculated by averaging over the distribution of infectious periods and the distribution of transmission rates [**43**]. This argument does not hold when, for instance, transmission rates and/or infectious periods are drawn from different distributions for different edges and/or nodes and the resulting transmission probabilities differ between edges. Newman [**43**] shows that the percolation approach can still be useful even in some such instances, although the problem is no longer described in terms of a single transmissibility.

[6]There are also clear analogies between the preceding discussions of the component size distribution of the random graph and the occurrence of minor and major outbreaks in the stochastic well-mixed model.

For the simple SIR processes introduced at the start of section 3, the transmissibility depends on the transmission rate along edges, β, and the infectious period distribution. If the duration of infection, τ, is the same for all individuals, we have that $T = 1 - \exp(-\beta\tau)$ [**28, 43**]. For the constant recovery model (exponentially distributed infectious periods with mean τ), averaging over the infectious period distribution gives $T = \beta\tau/(1 + \beta\tau)$ [**28**].

We notice that if the product $\beta\tau$ is small compared to one, then either of the above expressions for T can be well approximated by $\beta\tau$. Intuitively, this makes sense since the product of the per-link transmission rate and the average duration of infection will approximately equal the transmission probability along the link, provided that this product is not too large. (That $\beta\tau$ is just an approximation to the transmission probability is immediately clear since it can take values greater than one.)

3.2.1. *Percolation Theory and Epidemic Thresholds.* The results of percolation theory show that threshold behavior is typical in these systems. As in the earlier discussion of the construction of the random graph, there is a critical value of the traversal probability, below which the traversable network consists of a large number of small components, which have typical size $O(1)$, and above which the traversable network consists of a single giant component, of size $O(N)$, and a number of small components. In the epidemiological interpretation, there is a critical value of the transmissibility, below which only minor outbreaks occur and above which either major or minor outbreaks can occur. A major outbreak will occur if the initial infective is located in the giant component, and a minor outbreak will occur if the initial infective is located in one of the small components. The probability of the occurrence of a major outbreak, therefore, is equal to the fraction of nodes that are found within the giant component. This threshold behavior is precisely that of the epidemiologically familiar $R_0 = 1$ threshold, as described above in the case of the well-mixed stochastic model.

Percolation theory can be used to quantify the threshold and give a detailed account of behavior near the threshold (so-called 'critical behavior'), in addition to yielding information about the probability distribution of outbreak sizes and the probability of the occurrence of an outbreak. For example, Grassberger [**22**] describes power law (scaling) behavior in the outbreak probability when the transmissibility is just above the threshold, and in the outbreak size distribution when the transmissibility is just below threshold (see also [**54**]).

One important difference between network settings and the population-level well-mixed models is that individuals have a fixed set of contacts. This leads to some important differences in the expression for the basic reproductive number, even for the random graphs that attempt to model well-mixed settings. If, in a large well-mixed population (i.e. a random network), every individual is assumed to have exactly k neighbors, and it is assumed that there are no loops of short length in the network, then the basic reproductive number is given by [**15, 16**]

$$R_0 = T(k-1). \tag{8}$$

In the language of percolation theory, there is a critical value of the transmissibility, given by $T_C = 1/(k-1)$, above which large outbreaks can occur [**43**].

Notice that R_0 depends on the number of neighbors minus one. Every individual who transmits infection, except for the initial infective, can have at most $k-1$ susceptible neighbors because they must have acquired infection from one

of their neighbors. Also notice that the value of the basic reproductive number now depends on the infectious period distribution, although this dependence will be weak if $\beta\tau$ is small.

It is instructive to compare expression (8) to the simple-minded expression $R_0 = \beta k \tau$ obtained using the population-level well-mixed model. Two differences can be seen: the latter involves k and not $k-1$, and, since $\beta\tau$ can be greater than one, the average number of secondary infections can exceed k. These deficiencies arise because the well-mixed model does not account for individuals having a finite number of neighbors.

Using branching process theory, Diekmann and co-workers [**15, 16**] derived a formula, valid for $N \to \infty$, for the average final size of a major outbreak in the k-neighbor random graph described above, following the introduction of a single infective. If R_0 is greater than one, it can be shown that the following equation

$$\theta = (1 - T + \theta T)^{k-1} \tag{9}$$

has a unique solution for θ in $(0, 1)$. Diekmann *et al.* then showed that the final size of the epidemic (i.e. the fraction ever infected), y_∞, is given by

$$y_\infty = 1 - (1 - T + \theta T)^k . \tag{10}$$

As discussed earlier, y_∞ also gives the probability of the occurrence of a major outbreak. In the limit as $k \to \infty$, if β is scaled inversely with the connectivity[7], Diekmann *et al.* note that equation (6) is recovered.

Clique structure also impacts the spread of infection [**9, 26, 45**], since it reduces the number of secondary infections that each individual can cause. Using figure 2c as an example, imagine that individual A is the initial infected individual and that they infect person B. Although both individuals still have seven susceptible neighbors, four of these neighbors are shared and so the maximum number of further secondary infections is just ten. From the viewpoint of transmission, localized interactions mean that there are a large number of wasted contacts. Consequently, the value of the basic reproductive number is lower than in comparable well-mixed settings. This leads to outbreak sizes being smaller in cliqueish networks [**26, 45**], although Newman points out the counterintuitive result that cliqueishness can make it easier for these smaller outbreaks to occur when the transmissibility is low [**45**].

Threshold values of the transmissibility for lattices and small-world networks can can be obtained using percolation results. As mentioned above, percolation on regular lattices has received much attention and so critical values are known for many bond percolation problems. For the two dimensional lattice with the von Neumann neighborhood (i.e. connections are made to the four nearest neighbors), it is known that the percolation threshold occurs at $q = 1/2$ [**23**]. The percolation problem on small-world networks has also received much attention [**38, 39, 42, 47**]. Moore and Newman [**38, 39**] provide solutions for the bond percolation problems on one dimensional small-world networks for which connections are made either to the nearest neighbors or the two nearest neighbors on either side of each individual. They point out that their approach could be extended to more highly connected settings, but that the calculations rapidly become more involved as k increases. It

[7]This corresponds to the notion that as a person makes contacts with more people, they will spend less time in contact with each of them: their infectivity is diluted across their neighbors. This corresponds to the formula $cb = k\beta$ discussed in Section 3.1.

should also be pointed out that they describe their approach for the variant small-world generating algorithm in which long-range links are added in addition to those of the lattice, rather than the rewiring approach that we adopt here.

Heterogeneity of the network's degree distribution impacts the spread of infection. A well-known formula, appropriate for a particular mixing pattern known as proportionate mixing, illustrates how heterogeneity inflates the basic reproductive number [**2, 15, 36, 43**]:

$$R_0 = T\left(\langle k\rangle - 1 + \frac{\mathrm{Var}(k)}{\langle k\rangle}\right). \tag{11}$$

Here $\langle k\rangle$ and $\mathrm{Var}(k)$ denote the mean and variance of the degree distribution, respectively. Many real-world networks have highly heterogeneous degree distributions: in such instances, the expression for the basic reproductive number may be dominated by the variation in the degree distribution, rather than its average.

Andersson [**2**] (see also [**15, 36, 43**]) derived the following expressions that can be used to calculate the probability of the occurrence of a major outbreak and the average size of the resulting outbreak in this heterogeneous setting

$$\theta = \sum_{k=1}^{\infty} \frac{k\mu_k}{\langle k\rangle} \left(1 - T + \theta T\right)^{k-1} \tag{12}$$

$$y_\infty = 1 - \sum_{k=1}^{\infty} \mu_k \left(1 - T + \theta T\right)^{k}. \tag{13}$$

Here μ_k are the elements of the connectivity distribution. These expressions are analogous to equations (9) and (10), but are weighted by the connectivity distribution or, in the case of equation (12), by the distribution that depicts the probability that a randomly chosen contact is with an individual of connectivity k.

For the random graph, with its Poisson distributed degree distribution, equation (11) reduces to $R_0 = T\langle k\rangle$. This small adjustment echoes the earlier observation that the networks under consideration here are fairly homogeneous: the issue of heterogeneity is not a major concern for this study. (We remark that the R_0 formula looks like the one obtained for the well-mixed deterministic model in section 3.1. That expression, however, ignored heterogeneity and so is not directly comparable: incorporation of heterogeneity within the deterministic framework would also lead to an inflation of the value of R_0, and in a way that is quite analogous to equation (11) [**33, 51**].)

3.3. Pair Models. Percolation theory is not the only approach that has been used to address questions such as whether an infection can spread on a given network, the resulting outbreak size distribution or the temporal dynamics of this spread.

Infection spreads less rapidly on a typical network than it would in a well-mixed population. In the latter, a single infective can directly infect any susceptible individual because everyone interacts with everyone else in the population. Unless the network is completely connected, each individual will have fewer than $N-1$ neighbors, and so infection must travel via intermediate individuals in order to reach everyone in the population.

In contrast to well-mixed settings, transmission rates for general networks cannot be adequately described just in terms of the numbers of susceptibles and infectives. Rates of transmission depend on the configuration of these susceptibles and infectives, depending on the number of instances in which a susceptible is found to be in contact with an infective individual: we might term these instances susceptible-infective pairs.

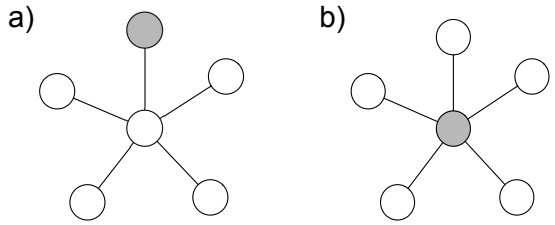

FIGURE 4. Spread of infection depends on the configuration of pairs. The network fragments in both (a) and (b) have five susceptibles (unshaded circles) and one infective (shaded circle). Infection will clearly spread more rapidly in (b) than in (a) because the single infective is connected to more susceptibles. In (a), we have four S-S (susceptible-susceptible) pairs and one S-I (susceptible-infective) pair, whereas in (b) we have five S-I pairs.

Pair models [**9, 19, 26, 53, 56**] attempt to capture this structure, keeping track of the numbers of pairs in which individuals of different types are connected. Differential equations then describe how the numbers of the different types of pairs change over time. Keeling [**26**] derives the following set of equations for the dynamics of pairs for the SIR process

$$[\dot{SS}] = -2\beta[SSI] \tag{14}$$

$$[\dot{SI}] = \beta\left([SSI] - [ISI] - [SI]\right) - \gamma[SI] \tag{15}$$

$$[\dot{SR}] = -\beta[RSI] + \gamma[SI] \tag{16}$$

$$[\dot{II}] = 2\beta\left([ISI] + [SI]\right) - 2\gamma[II] \tag{17}$$

$$[\dot{IR}] = \beta[RSI] + \gamma\left([II] - [IR]\right). \tag{18}$$

Here, the number of X-Y pairs is denoted by $[XY]$, where X and Y can denote S, I or R type individuals. For book-keeping reasons, $[XX]$ denotes twice the number of X-X pairs. The number of X-Y-Z triples is written as $[XYZ]$. Notice that it is not necessary to have equations that track the numbers of individuals of each type, which in Keeling's notation would be denoted $[X]$: these quantities can always be calculated in terms of the numbers of pairs involving the type of interest. For example, the number of susceptible individuals can be calculated if the numbers of S-S, S-I and S-R pairs are known.

One feature of these equations is that the rates of change of pairs involve the configuration of triples. If one looked at equations for triples, one would find that they involve quartets, and so on. In order to usefully employ this approach, the set of equations is truncated (or closed) at some order. One way to do this is via so-called pair approximations, in which the configuration of triples is described, in

an approximate way, in terms of the configuration of pairs. This then eliminates the need for equations that describe how the numbers of triples evolve over time, leaving a closed set of equations for the numbers of pairs. We remark that the use of the pair approach for this SIR process replaces a two dimensional model, equations (2) and (3) by a five dimensional system.

Various pair approximations can be made, reflecting the geometry of the network. A different approximation would be more suited to graphs that are similar to random graphs than would be appropriate for ones that are similar to lattices. Keeling [**26**] considers graphs for which each individual has k neighbors and that are further specified in terms of their cliqueishness, as measured by the transitivity, ϕ. He uses the following approximation

$$[ABC] \approx \frac{k-1}{k}\frac{[AB][BC]}{[B]}\left((1-\phi)+\phi\frac{N}{k}\frac{[AC]}{[A][C]}\right), \tag{19}$$

first described by Morris [**40**], to close the set of moment equations, where N is the population size. Bauch [**9**] calls this the triangular pair approximation, since it accounts for the occurrence of triangles in the network. The accuracies of various pair approximations in lattices with various different geometries are discussed by Bauch [**9**], in which it is pointed out that the standard pair approximation often does not provide a completely satisfactory description. (For this reason, the predictions made by pair models are almost invariably compared to those obtained by numerical simulation of the full network model.)

A number of techniques are available that extend these pair models [**8, 9, 19**], including extending the set of equations to account for the configuration of triples [**8, 9**]. Alternatively, in a disease invasion setting, the so-called invasory pair approximation attempts to account for the structure of the clusters that develop during the invasion process [**9**]. Another approach for invasion settings, the pair-edge approximation [**19**], attempts to describe the leading edge of the invasion wave front, enabling estimation of the speed at which the infection spreads across the population.

Threshold results can be obtained using pair models, and their extensions, by examining whether infection can invade a population or not. This question is more difficult to address in the five dimensional pair model than it is in the standard deterministic SIR model. This task can be simplified somewhat on account of the spatial structure that quickly develops during the invasion of the infection. This allows quasi steady state assumptions to be made, reducing the dimensionality of the system and facilitating the calculation of thresholds [**9, 19, 26**].

The system of pair equations can be integrated, allowing the final size of the epidemic to be calculated. Typically this integration can only be carried out numerically.

4. Epidemic Dynamics

The impact of various aspects of network structure on the spread of infection is easily understood in terms of the preceding discussion. All other things being equal, infection will spread more readily and rapidly when the average connectivity of individuals is higher. Highly localized mixing hinders the spread of infection, both because of cliqueishness, which leads to wasted contacts, and because long path

lengths mean that infection must travel through a large number of intermediates in order to spread across the entire network.

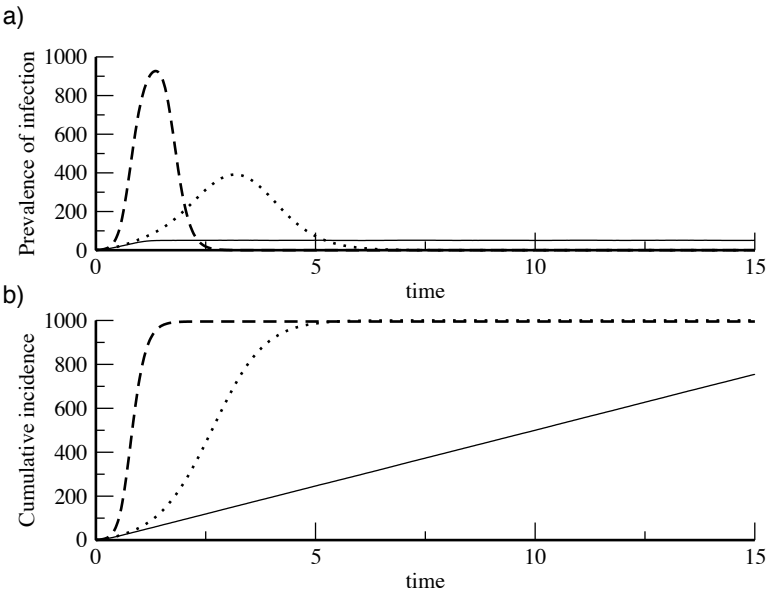

FIGURE 5. Averaged timecourse of epidemics on a few of the networks depicted in figure 3, generated using the Watts-Strogatz method, under the SIR infection process assuming that $\beta = 1.0$ and that each infection lasts exactly one time unit (i.e. $\tau = 1$). The upper panel depicts the prevalence of infection (i.e. the number of infectious people at each point in time) and the lower panel depicts the cumulative incidence (i.e. the total number of cases up to the time point). The solid curves illustrate simulations on the lattice ($p = 0$), and the dashed curves illustrate simulations on the random graph generated by rewiring all the links of the lattice ($p = 1$). The dotted curves denote the intermediate case of a small-world network, with just one percent of the lattice's links being rewired ($p = 0.01$). Curves are calculating by averaging over 1000 realizations of the model.

Epidemics on lattices, therefore, spread slowly and exhibit high degrees of spatial structure. If infection is introduced at a single location, this spread will take the form of an outward spreading wave, the speed of which will depend on epidemiological parameters and the geometry of the lattice. Epidemics on random graphs spread much more rapidly and exhibit little or no spatial structure. Epidemics on small-world networks will be intermediate between these two extremes, and so spread at some intermediate speed. Their spatial structure will involve a number of growing clusters of infection: the mainly local nature of mixing means that there will be wave-like spread out from a point of introduction, but long-range transmission events will often take infection to virgin territory, giving rise to new clusters of infection.

Figures 5-7 illustrate several aspects of epidemic dynamics, summarizing the results of repeated stochastic simulation of the SIR process on graphs generated by the Watts-Strogatz small-world algorithm. Figure 5 depicts the timecourse of epidemics, averaged over 1000 realizations of the model, assuming three different values of the Watts-Strogatz per-link rewiring probability. In this figure, the parameter values are chosen so that the transmissibility is quite some way above the epidemic threshold. The spread of infection is much faster in the random graph than in the lattice. Following a short initial transient, in this case so short that it is barely noticeable on the scale of the figure, the cumulative incidence increases linearly with time for a long period in the lattice, reflecting the wave-like spread of infection outwards from the point of introduction. The impact of long-range links is dramatic: a rewiring probability of just one percent turns the lattice into a network on which infection spreads considerably faster.

An interesting observation [**30**] is that the timecourse of the epidemic in the small-world regime echoes that of a standard SIR model, albeit with different values of its parameters. This idea has recently been taken up by Aparacio and Pascual (manuscript in prep.) who modify a standard SIR model to capture some of the features of spatial structure, such as the local depletion of susceptibles within infection clusters, providing a description of the epidemic in terms of a low dimensional dynamical system.

Figure 6 illustrates the distribution of outbreak sizes that are seen when a single infective is introduced into an otherwise susceptible population. The parameter values in figure 6 are chosen to be closer to the epidemic threshold than those of the previous figure, with the transmissibility of this infection equalling $T = 1 - \exp(-0.2) \approx 0.181$. For the completely rewired network, use of expression (8), derived for the k-neighbor random graph setting, gives $R_0 \approx 1.63$. Accounting for the heterogeneity of the random graph, expression (11) gives the slightly higher value of $R_0 \approx 1.81$. In either case, the R_0 value (for the totally rewired network) is not so far above the epidemic threshold.

For the parameter values in figure 6, rewiring the lattice has moved the population from being below the epidemic threshold to above the epidemic threshold. For the lattice (figure 6a), we see that introduction of infection only leads to a minor outbreak, whereas for the corresponding small-world network (if p is sufficiently large), or the random graph obtained when all links are rewired, introduction either leads to a minor outbreak or a major outbreak. The histograms in figures 6c and 6d are clearly bimodal, comprising two components: the distribution of minor outbreaks and the distribution of major outbreaks.

As in our earlier discussion of the stochastic well-mixed model, we see that, even above the epidemic threshold, introduction of infection does not guarantee the occurrence of a major outbreak. Looking at the case of the totally rewired network, we see that a major outbreak is seen in about 74% of the simulations shown in figure 6d (see also figure 7, panel b). Approximating the rewired network by a k-neighbor random graph, equations (9) and (10) give the probability of the occurrence of a major outbreak, following the introduction of a single infective, as ≈ 0.749. Taking the degree distribution of the totally rewired network to be Poisson, and accounting for the heterogeneity of the totally rewired network, formulae (12) and (13) give the probability of the occurrence of a major outbreak as ≈ 0.737. These predicted values are not so different from each other (on account of the

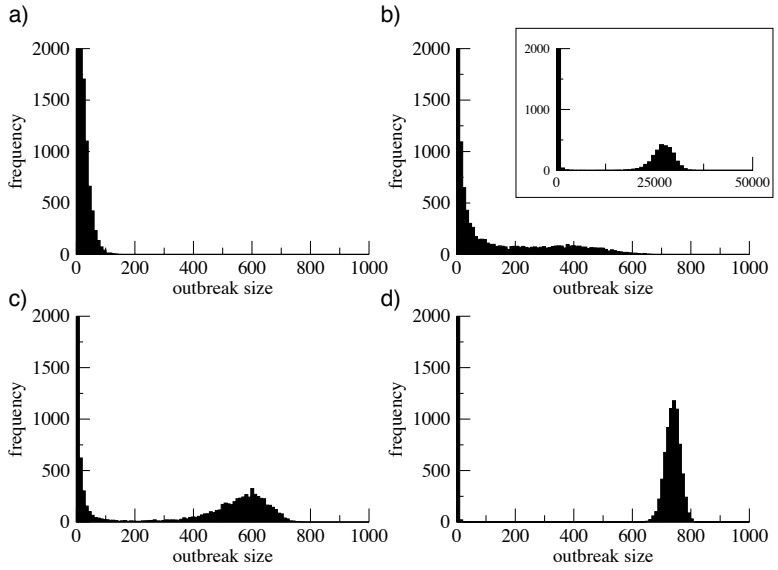

FIGURE 6. Outbreak size distributions for the SIR infection process on the networks of figure 3, assuming that $\beta = 0.2$ and that each infection lasts exactly one time unit ($\tau = 1$). In each case, the average connectivity of the network is ten and the network consists of 1000 individuals. In panel (a), no links are rewired. Panels (b) and (c) use small-world networks generated from the underlying 1D lattice with per-link rewiring probabilities equal to 0.040 and 0.079. All links are rewired in the network of panel (d). The inset in panel (b) shows the corresponding outbreak size distribution when the network contains 100 000 individuals (with p still equal to 0.040). Each panel depicts the outcomes of 10000 realizations of the model, but for clarity the vertical axis is cut off at 2000. (In each case, the peak corresponding to ten or fewer cases reaches above this cut-off.)

relatively low heterogeneity of the network's degree distribution) and are both in good agreement with the simulation results.

From the histograms of figure 6, we see that increasing the rewiring probability increases the probability of the occurrence of a major epidemic and increases the average size of the ensuing outbreak. We also notice that the distribution of the sizes of the major epidemics becomes less variable as the rewiring probability increases. Close to the threshold, we see a transition regime in which the distinction between minor and major outbreaks is less clear: the histogram in figure 6b is less clearly bimodal than one obtained further above the threshold (for example, figure 6c). Since the minor outbreaks are of typical size $O(1)$ and the major outbreaks are of size $O(N)$, this transition regime is less noticeable when N is larger. This is illustrated in the inset to panel (b), obtained for the same value of p but with a population size that is one hundred times larger. Comparison between the main graph of panel (b) and its inset also illustrates that the sizes of the major outbreaks

are more closely centered around their mean for large population sizes (see [**3**] for a more complete discussion of this effect).

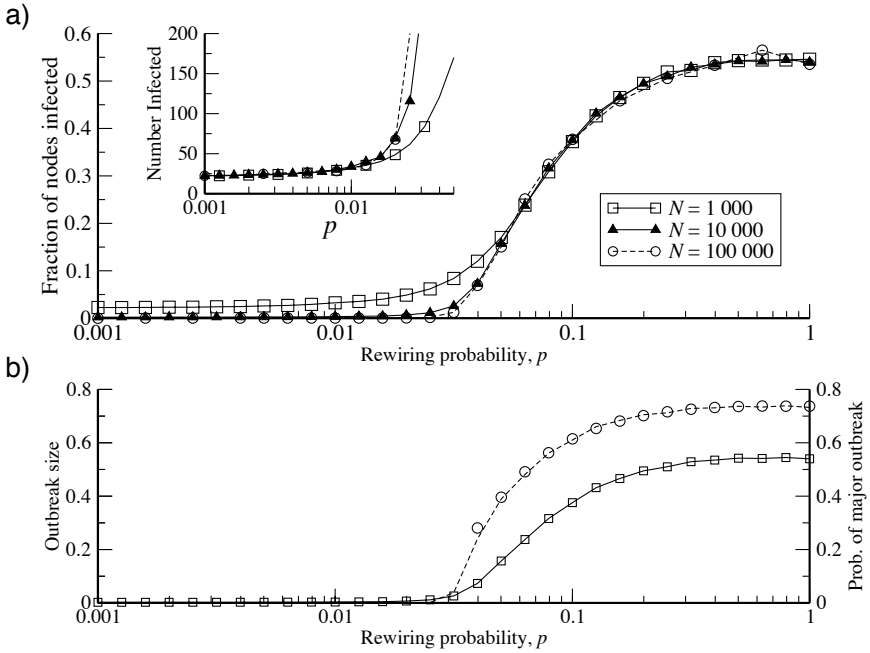

FIGURE 7. Panel (a): Dependence of the average **fraction** of nodes that ever become infected (main graph) and the average **number** of nodes that ever become infected (inset graph), following the introduction of a single infective individual, on the rewiring probability in the Watts-Strogatz small-world network algorithm. Networks are generated as in figure 3, with population sizes of $N = 1000$ (solid line with squares), $N = 10000$ (solid line with filled triangles) or $N = 100000$ (dashed line with circles). Infection parameters are $\beta = 0.2$ and $\tau = 1$. Averages were taken over 10 000 realizations for $N = 1000$ and $N = 10000$, but over 1000 realizations for $N = 100000$. Panel (b) depicts, in terms of the rewiring probability, three quantities: the average outbreak size (solid curve with squares), the probability of the occurrence of a major outbreak (dashed curve) and the average size of a major outbreak, if one occurs (circles). All curves in panel (b) are calculated using the simulations shown in panel (a) for $N = 10000$. Notice that the 'average size' curve is calculated by averaging over all realizations of the model (regardless of whether they involved a minor or a major outbreak), and thus corresponds to one of the curves in panel (a).

Figure 7a clearly illustrates the threshold behavior of the model as the rewiring probability is increased. Also visible is the $O(N)$ behavior of the sizes of the major outbreaks, with the average fraction of the population ever infected above

threshold being independent of N (main graph). The $O(1)$ behavior of the sizes of minor outbreaks is also clear, with the average number of people ever infected below threshold being independent of N (inset graph).

Figure 7b depicts, in terms of the rewiring parameter, the average outbreak size (calculated across all realizations, regardless of whether a minor or major outbreak occurred), the probability of the occurrence of a major outbreak and the size of the ensuing major outbreak, if it occurs. (Put another way, the last quantity is the average outbreak size conditional on the occurrence of a major outbreak.) As expected, the probability of the occurrence of a major outbreak is, within the accuracy of the simulations, equal to the (fractional) size of a major outbreak for values of p above threshold. The transition region around the threshold imposes some difficulties in the estimation of the probability and size of major outbreaks in figure 7b. Since threshold behavior is sharper for larger values of N, as explained in the discussion of figure 6b and its inset, these difficulties are mitigated by employing larger population sizes (hence the use of $N = 10000$ in figure 7b rather than the $N = 1000$ of figure 6).

5. Dynamics in Endemic Settings

The assumptions of lifelong immunity and the neglect of demographic processes are unrealistic for many infections. In many instances, replenishment of the susceptible population, either with the waning of immunity or the birth of new susceptibles, has a major impact on the dynamics of the infection. Most importantly, this process can lead to the establishment of an equilibrium level of infection in which the rate at which susceptibles are infected (or die) balances the rate at which susceptibles are recruited and the rate at which individuals become infected balances the rate at which they recover (or die) [**1, 15**].

Modeling demographic processes requires the removal of individuals from the network upon their death and the insertion of individuals as they are born. Unless one simply replaces a dying individual by a newborn, this raises questions as to where newborns should be placed in the network. From the viewpoint of network settings, therefore, it is much easier to model the replenishment of susceptibles by allowing immunity to wane. The SIRS model assumes that recovered individuals lose their immunity at some point. A limiting case of the SIRS model is the SIS model, in which the duration of immunity is simply set equal to zero: individuals recover to a susceptible state. On the other hand, if the duration of immunity is allowed to tend to infinity, individuals never recover and the SIRS model approaches the SIR model.

For simple well-mixed models, there are many similarities between the behavior of SIR and either SIS or SIRS infection processes. Most notably, the threshold concept carries over to these processes, with the basic reproductive number determining a disease invasion condition: introduction of infection can lead to an outbreak when the basic reproductive number is greater than one. In addition, the same condition ensures the existence and stability of an endemic equilibrium (i.e. one with a positive prevalence of infection) of the system [**1, 15, 44**]. In other words, the $R_0 = 1$ condition is also a disease persistence threshold. For many models, the expression for R_0 is identical in both epidemic (e.g. SIR) and endemic (e.g. SIRS) settings.

In non well-mixed models, the invasion and persistence thresholds need not be identical: one scenario under which this can arise involves the occurrence of a so-called "backwards bifurcation" at $R_0 = 1$, and has been the subject of many studies [**12, 18**].

The stability of the endemic equilibrium has been of considerable interest: in simple settings, the existence of a (positive) endemic equilibrium guarantees its stability. This is not always the case, however, and for many models, the endemic equilibrium can become unstable via a Hopf bifurcation, leading to stable oscillations in the prevalence of infection [**25**].

In stochastic settings, random fluctuations prevent the system from settling into an equilibrium. If these fluctuations are large enough, then the infection can go extinct: if the number of infectives falls close to zero, then it is possible for all of them to recover before passing on the infection. This phenomenon is known as (endemic) fade-out [**7**]: stochasticity means that an infection may not persist even though its basic reproductive number is large compared to one. This is an important aspect of the dynamics of many infections, such as measles, in all but the largest-sized cities. In spatial settings, fade-out leads to local extinction of infection, but need not link to global extinction of infection if outbreaks in different locations are desynchronized [**10, 41**]. Indeed, movement of infectious individuals means that infection can be reintroduced into a region where it previously underwent fade-out. Consequently, the population-level persistence of infection depends in an important way on the synchrony of outbreaks across the population [**10, 27**].

Much of the discussion of the impact of spatial structure on the dynamics and persistence of infection has been framed in terms of metapopulation models. These questions have, however, also been investigated within the small-world framework [**31, 58**].

Kuperman and Abramson [**31**] investigated the dynamics of an SIRS process on a small-world network. For small values of the Watts-Strogatz rewiring parameter, p, they found that the population-level prevalence exhibited small stochastic fluctuations around an endemic equilibrium. Quite different behavior was seen for larger values of the rewiring parameter, however, with the prevalence undergoing large amplitude oscillations. These oscillations occurred on the timescale of the duration of immunity. The change in behavior was accompanied by a change in the synchrony of outbreaks between different regions of their population.

An intuitive explanation of these results can be given in terms of the mixing pattern of the population and the resulting level of synchrony. For small values of p, the mixing is mainly local, and so oscillatory behaviors in different locales occur in an asynchronous fashion. With little synchrony between outbreaks in different regions of the population, the population level prevalence undergoes small amplitude fluctuations about some level: the local oscillations are averaged out at the level of the whole population (see also [**17**]). When p is large, mixing is global in nature and so there is a greater tendency for synchrony between outbreaks in different parts of the population. Indeed, the oscillations in prevalence are predicted by a mean-field (well-mixed) model for the system [**25**].

Verdasca and co-workers [**58**] considered the persistence and synchrony of outbreaks of an infection such as measles on small-world networks using what they described as an SIR process with demography. (Actually, their model employed an

SIRS process since they modeled births by replacing recovered individuals by susceptibles.) Spatial correlations were found to enhance the stochastic fluctuations around the endemic equilibrium, with smaller fluctuations in the well-mixed case ($p = 1$) than when the network was close to a lattice ($p \approx 0$).

As noted by Verdasca et al., their study and that of Kuperman and Abramson make apparently conflicting predictions regarding the impact of spatial structure on oscillatory behavior. Although some aspects of this conflict are not fully resolved, Verdasca et al. do point out that the two studies examine quite different parameter regimes. Kuperman and Abramson's study assumes that immunity lasts for a time that is comparable to the duration of infection; the resulting oscillations are on a similar timescale. For Verdasca et al., immunity is lifelong, while infection lasts for at most a few weeks; the resulting oscillations are on an intermediate timescale, being on the order of one to three years. Kuperman and Abramson point out that oscillations are not seen in their model when there is a marked difference between the durations of infectiousness and immunity.

The conventional wisdom in the literature on the metapopulation dynamics of measles is that spatial structure enhances the persistence of infection [**10, 27**] since outbreak asynchrony can allow global persistence even in the face of frequent local extinction. Persistence is much less likely in a well-mixed population than in a more structured population. Verdasca et al., however, described an abrupt transition in the persistence of infection as the rewiring probability, p, was increased. In their small-world model, persistence was much more likely to occur in networks that were close to well-mixed. They described this transition in terms of a percolation threshold and so it appears to be related to the increase in R_0 that accompanies higher values of p, as described in the previous sections. Since it is well-known that the invasion threshold is exceeded for measles— indeed with an R_0 of about 15, measles is highly infectious— this observation of Verdasca et al. does not, unfortunately, address the main question of interest for measles, namely the impact of spatial structure in countering the endemic fade-out effect.

Questions of dynamics and persistence have been considered in other network settings; these results shed some light on the above studies of small-world settings. Morris [**40**], discussed by Rand [**53**], found that the stable endemic equilibrium of the SIR model gave way to limit cycle behavior on clustered networks as the quantity ϕ increased through a critical value. We remark that the direction in which this change occurs is in agreement with the results of Verdasca et al. Keeling and co-workers [**29**] demonstrated that clustering in a network could enhance the persistence of an endemic infection, a finding that is in agreement with the results of metapopulation approaches.

6. Discussion

Network models provide a flexible framework within which the impact of population structure on the transmission dynamics of infections can be studied. On one level, they provide a research tool that can be used to investigate the relationship between epidemiological interactions that are described in terms of individuals and their behavior, and the resulting population-level patterns of disease transmission. In this way, network models can be used to better understand the assumptions that underlie many population-level epidemiological models. On another level, network models are being increasingly used as a practical tool, to aid in the analysis

of epidemiological data (see [**35**]), to provide ever more realistic models for real-world populations (see [**20**]) and to aid in the design of control strategies (see [**21, 35, 36, 52**])

Although we have framed our discussion within the context of the spatial structure of populations, the network need not correspond to geographic space but rather to some more general notion of the social space of a population. For instance, mixing might be age dependent: for instance, children of the same age are more likely to spend time together at school [**29**].

Heterogeneity in the connectivity distribution is an important aspect of the structure of many populations, but is something we have not discussed here in detail. As mentioned above, the degree distributions of the Watts-Strogatz small-world networks exhibit only low levels of heterogeneity: this class of networks is, therefore, a poor model for many real world populations. It has long been realized, for instance, that sexual partnership networks exhibit extreme heterogeneity (see, for example, [**1**]), that this heterogeneity has a major impact on the spread of infection and that models that ignore this heterogeneity will provide poor descriptions of reality.

Much of the discussion of network models has been within the context of disease invasion. This setting has long been favored by modelers since invasion is the setting within which mathematical analysis of epidemic models is often easiest. More recently, real world events have also focused attention on disease invasion: concerns about bioterrorism have fueled the development of models that describe the deliberate introduction of infectious diseases into urban and other settings. Network approaches have provided a natural framework for such models [**13, 20**]. For many infections, however, the most interesting epidemiological questions concern endemic settings. Although there is an ever growing literature concerning network approaches for endemic infections, much work remains to be done here.

7. Acknowledgements

We wish to thank several anonymous referees for their insightful comments, which greatly helped in improving this paper.

References

[1] R. M. Anderson and R. M. May, *Infectious diseases of humans: Dynamics and control*, Oxford University Press, Oxford, 1991.

[2] H. Andersson, *Limit theorems for a random graph epidemic model*, Ann. Appl. Prob. **8** (1998), 1331–1349.

[3] H. Andersson and T. Britton, *Stochastic epidemic models and their analysis*, Springer-Verlag, 2000.

[4] F. G. Ball, D. Mollison, and G. Scalia-Tomba, *Epidemics with two levels of mixing*, Ann. Appl. Prob. **7** (1997), 46–89.

[5] A. Barbour and D. Mollison, *Epidemics and random graphs*, Stochastic Processes in Epidemic Theory (J.-P. Gabriel, C. Lefévre, and P. Picard, eds.), Lecture Notes in Biomathematics, vol. 86, Springer-Verlag, 1990, pp. 86–89.

[6] A. Barrat and M. Weigt, *On the properties of small-world network models*, Eur. Phys. J. B **13** (2000), 547–560.

[7] M. S. Bartlett, *Deterministic and stochastic models for recurrent epidemics*, Proc. 3rd Berkeley Symp. on Math., Stat., and Prob. **4** (1956), 81–109.

[8] C. T. Bauch, *A versatile ODE approximation to a network model for the spread of sexually transmitted diseases*, J. Math. Biol. **45** (2002), 375–395.

[9] ———, *The spread of infectious diseases in spatially structured populations: An invasory pair approximation*, Math. Biosci. **198** (2005), 217–237.
[10] B. M. Bolker and B. T. Grenfell, *Impact of vaccination on the spatial correlation and persistence of measles dynamics*, Proc. Natl. Acad. Sci. USA **93** (1996), no. 22, 12648–12653.
[11] B. Bollobás, *Random graphs*, Academic Press, 1985.
[12] C. Castillo-Chavez, K. Cooke, W. Huang, and S. A. Levin, *On the role of long incubation periods in the dynamics of HIV/AIDS, Part 2: Multiple group models*, Mathematical and Statistical Approaches to AIDS Epidemiology (C. Castillo-Chavez, ed.), Lecture Notes in Biomathematics, vol. 83, Springer-Verlag, 1989, pp. 200–217.
[13] G. Chowell and C. Castillo-Chávez, *Worst-case scenarios and epidemics*, Bioterrorism: Mathematical Modeling Applications in Homeland Security (H. T. Banks and C. Castillo-Chávez, eds.), SIAM Frontiers in Applied Mathematics, vol. 28, SIAM, 2003.
[14] F. Chung and L. Lu, *The average distances in random graphs with given expected degrees*, Proc. Natl. Acad. Sci. USA **99** (2002), 15879–15882.
[15] O. Diekmann and J. A. P. Heesterbeek, *Mathematical epidemiology of infectious diseases*, John Wiley & Son, Chichester, 2000.
[16] O. Diekmann, M. C. M. De Jong, and J. A. J. Metz, *A deterministic epidemic model taking account of repeated contacts between the same individuals*, J. Appl. Prob. **35** (1998), 448–62.
[17] R. Durrett and S. A. Levin, *The importance of being discrete (and spatial)*, Theor. Popul. Biol. **46** (1994), 363–394.
[18] J. Dushoff, *Incorporating immunological ideas in epidemiological models*, J. Theor. Biol. **180** (1996), no. 3, 181–187.
[19] S. P. Ellner, A. Sasaki, Y. Haraguchi, and H. Matsuda, *Speed of invasion in lattice population models: pair-edge approximation*, J. Math. Biol. **36** (1998), 469–484.
[20] 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 settings*, Nature **429** (2004), 180–184.
[21] N. M. Ferguson, D. A. Cummings, S. Cauchemez, C. Fraser, S. Riley, A. Meeyai, S. Iamsirithaworn, and D. S. Burke, *Strategies for containing an emerging influenza pandemic in southeast asia*, Nature **437** (2005), 209–214.
[22] P. Grassberger, *On the critical behavior of the general epidemic process and dynamical percolation*, Math. Biosci. **63** (1983), 157–172.
[23] G. Grimmett, *Percolation*, 2nd edn., Springer-Verlag, 1991.
[24] T. E. Harris, *Contact interactions on a lattice*, Ann. Prob. **2** (1974), 969–988.
[25] H. W. Hethcote, H. W. Stech, and P. Van Den Driessche, *Nonlinear oscillations in epidemic models*, SIAM J. Appl. Math. **40** (1981), 1–9.
[26] M. J. Keeling, *The effects of local spatial structure on epidemiological invasions*, Proc. R. Soc. Lond. B. **266** (1999), 859–867.
[27] ———, *Multiplicative moments and measures of persistence in ecology*, J. Theor. Biol. **205** (2000), 269–281.
[28] M. J. Keeling and B. T. Grenfell, *Individual-based perspectives on R(0)*, J. Theor. Biol. **203** (2000), 51–61.
[29] M. J. Keeling, D. A. Rand, and A. J. Morris, *Correlation models for childhood epidemics*, Proc. R. Soc. Lond. B **264** (1997), 1149–1156.
[30] A. Kleczkowski and B. T. Grenfell, *Mean-field-type equations for spread of epidemics: the 'small world' model*, Physica A **274** (1999), 355–360.
[31] M. Kuperman and G. Abramson, *Small world effect in an epidemiological model*, Phys. Rev. Lett. **86** (2001), 2909–2912.
[32] S. A. Levin and R. Durrett, *From individuals to epidemics*, Phil. Trans. Roy. Soc. Lond. B **351** (1996), 1615–1621.
[33] R. M. May and A. L. Lloyd, *Infection dynamics on scale-free networks*, Phys. Rev. E **64** (2001), 066112.
[34] J. A. J. Metz, D. Mollison, and F. van den Bosch, *The dynamics of invasion waves*, The Geometry of Ecological Interactions: Simplifying Spatial Complexity (U. Dieckmann, R. Law, and J. A. J. Metz, eds.), Cambridge University Press, Cambridge, 2000.
[35] L. A. Meyers, M. E. Newman, M. Martin, and S. Schrag, *Applying network theory to epidemics: control measures for Mycoplasma pneumoniae outbreaks*, Emerg. Infect. Dis. **9** (2003), 204–210.

[36] L. A. Meyers, B. Pourbohloul, M. E. Newman, D. M. Skowronski, and R. C. Brunham, *Network theory and SARS: predicting outbreak diversity*, J. Theor. Biol. **232** (2005), 71–81.
[37] D. Mollison, *Spatial contact models for ecological and epidemic spread*, J. R. Stat. Soc. B **39** (1977), 283–326.
[38] C. Moore and M. E. J. Newman, *Epidemics and percolation in small-world networks*, Phys. Rev. E **61** (2000), 5678–5682.
[39] C. Moore and M. E. J. Newman, *Exact solution of site and bond percolation on small-world networks*, Phys. Rev. E **62** (2000), 7059–7064.
[40] A. J. Morris, *Representing spatial interactions in simple ecological models*, Ph.D. thesis, Warwick University, 1997.
[41] G. D. Murray and A. D. Cliff, *A stochastic model for measles epidemics in a multi-region setting*, Trans. Inst. Brit. Geog. **2** (1975), 158–174.
[42] M. E. Newman, I. Jensen, and R. M. Ziff, *Percolation and epidemics in a two-dimensional small world*, Phys. Rev. E **65** (2002), 021904.
[43] M. E. J. Newman, *Spread of epidemic diseases on networks*, Phys. Rev. E **66** (2002), 016128.
[44] ———, *Ego-centered networks and the ripple effect*, Social Networks **25** (2003), 83–95.
[45] ———, *Properties of highly clustered networks*, Phys. Rev. E **68** (2003), 026121.
[46] ———, *The structure and function of complex networks*, SIAM Rev. **45** (2003), 167–256.
[47] M. E. J. Newman, C. Moore, and D. J. Watts, *Mean-field solution of the small-world network model*, Phys. Rev. Lett. **84** (2000), 3201–3204.
[48] M. E. J. Newman and D. J. Watts, *Scaling and percolation in the small-world network model*, Phys. Rev. E **60** (1999), 7332–7342.
[49] M. E. J. Newman, D. J. Watts, and S. H. Strogatz, *Random graph models of social networks*, Proc. Natl. Acad. Sci. USA **99** (2002), 2566–2572.
[50] J. V. Noble, *Geographic and temporal development of plague*, Nature **250** (1974), 726–728.
[51] R. Pastor-Satorras and A. Vespignani, *Epidemic spreading in scale-free networks*, Phys. Rev. Lett. **86** (2001), 3200–3203.
[52] B. Pourbohloul, L. A. Meyers, D. M. Skowronski, M. Krajden, D. M. Patrick, and R. C. Brunham, *Modeling control strategies of respiratory pathogens*, Emerg. Inf. Dis. **11** (2005), 1249–1256.
[53] D. A. Rand, *Correlation equations and pair approximations for spatial ecologies*, Advanced Ecological Theory: Principles and Applications (J. McGlade, ed.), Blackwell, 1999, pp. 100–142.
[54] C. J. Rhodes and R. M. Anderson, *Power laws governing epidemics in isolated populations*, Nature **381** (1996), 600–2.
[55] L. M. Sander, C. P. Warren, I. M. Sokolov, C. Simon C, and J. Koopman, *Percolation on heterogeneous networks as a model for epidemics*, Math. Biosci. **180** (2002), 293–305.
[56] K. Sato, H. Matsuda, and A. Sasaki, *Pathogen invasion and host extinction in lattice structured populations*, J. Math. Biol. **32** (1994), 251–268.
[57] F. van den Bosch, J. A. J. Metz, and O. Diekmann, *The velocity of spatial population expansion*, J. Math. Biol. **28** (1990), 529–565.
[58] J. Verdasca, M. M. Telo da Gama, A. Nunes A, N. R. Bernardino, J. M. Pacheco, and M. C. Gomes, *Recurrent epidemics in small world networks*, J. Theor. Biol. **233** (2005), 553–561.
[59] D. J. Watts and S. H. Strogatz, *Collective dynamics of 'small-world' networks*, Nature **393** (1998), 440–442.

Biomathematics Graduate Program and Department of Mathematics, North Carolina State University, Raleigh, NC 27695
E-mail address: alun_lloyd@ncsu.edu

Department of Epidemiology, University of North Carolina, School of Public Health, Chapel Hill, NC 27599
E-mail address: valeika@email.unc.edu

Center for Applied Mathematics, Cornell University 657 Rhodes Hall, Ithaca NY 14853
E-mail address: ariel@cam.cornell.edu