ELSEVIER

Mathematical Biosciences 166 (2000) 45–68

Mathematical Biosciences

www.elsevier.com/locate/mbs

Analysis and simulation of a stochastic, discrete-individual model of STD transmission with partnership concurrency

Stephen E. Chick [a,*], Andrew L. Adams [b], James S. Koopman [c]

[a] *Department of Industrial and Operations Engineering, The University of Michigan, 1205 Beal Avenue, Ann Arbor, MI 48109-2117, USA*

[b] *Merit Network Inc., 4251 Plymouth Road, Suite C, Ann Arbor, MI 48105-2785, USA*

[c] *Department of Epidemiology, The University of Michigan, 109 Observatory Street, Ann Arbor, MI 48109-2029, USA*

Received 22 November 1999; received in revised form 18 April 2000; accepted 11 May 2000

Abstract

Deterministic differential equation models indicate that partnership concurrency and non-homogeneous mixing patterns play an important role in the spread of sexually transmitted infections. Stochastic discrete-individual simulation studies arrive at similar conclusions, but from a very different modeling perspective. This paper presents a stochastic discrete-individual infection model that helps to unify these two approaches to infection modeling. The model allows for both partnership concurrency, as well as the infection, recovery, and reinfection of an individual from repeated contact with a partner, as occurs with many mucosal infections. The simplest form of the model is a network-valued Markov chain, where the network's nodes are individuals and arcs represent partnerships. Connections between the differential equation and discrete-individual approaches are constructed with large-population limits that approximate endemic levels and equilibrium probability distributions that describe partnership concurrency. A more general form of the discrete-individual model that allows for semi-Markovian dynamics and heterogeneous contact patterns is implemented in simulation software. Analytical and simulation results indicate that the basic reproduction number R_0 increases when reinfection is possible, and the epidemic rate of rise and endemic levels are not related by $1 - 1/R_0$, when partnerships are not point-time processes.

Keywords: Markovian infection process; Network-valued epidemic model; Pair formation model; Mucosal infection; Discrete-event simulation

* Corresponding author. Tel.: +1-734 763 2238; fax: +1-734 764 3451.
E-mail address: sechick@engin.umich.edu (S.E. Chick).

0025-5564/00/$ - see front matter
PII: S0025-5564(00)00028-6

1. Introduction

Infection levels are strongly determined by whether or not the contacts that transmit infection are made repeatedly with the same individual in an ongoing relationship, the potential for simultaneous ongoing relationships, and whether or not reinfection from a partner is possible [1–26]. One approach to studying these determinants uses continuous differential equations. Another is to study discrete-individual models with stochastic simulation. Until recently, practical and theoretical links between these two approaches have been lacking. GERMS [2] is a discrete-individual simulation model that has been proposed for analyzing transmission dynamics for sexually transmitted infections in social networks [16] that also allows for a continuous differential equation analysis in many situations. The purpose of this paper is to discuss the stochastic model behind GERMS, and to provide links between continuous differential equations and the discrete-individual simulations that describe special cases of that stochastic model.

The model structure of GERMS also allows for loss and reacquisition of infection in the context of a continuing relationship between individuals. This is probably quite common for mucosal infections, including gonorrhea and chlamydia. The analysis of the GERMS model structure below provides a new understanding of how endemic infection levels are determined when reinfection is a possibility. The analysis also extends the concept of basic reproduction number as defined for continuous and homogeneous population models to discrete-individual models and makes clear the limitations to extending this concept to situations where contacts continue rather than consisting of a point-time encounter. The analysis and interpretation are based on both theoretical and computer simulation results.

Section 2 examines a simple stochastic model of transmission for a susceptible–infectious–susceptible (SIS) infection in a homogeneously mixing population. In its simplest form, that stochastic discrete-individual SIS model can be studied as a deterministic compartmental model when a large-population approximation is made. Pseudo-equilibrium prevalence levels are derived using large population limit of a Markovian, closed population model when individuals are sequentially monogamous. The increase in the basic reproduction number R_0 is quantified when reinfection is possible, as may occur with mucosal infections. A Markov chain formulation indicates that R_0 is a more complicated concept when transmission occurs during ongoing partnerships, rather than point-time encounters. In particular, the familiar formula $1 - 1/R_0$ does *not* relate the endemic infection level to the expected number of secondary cases generated by a typical case in the early stages of an epidemic when all contacts are susceptible.

Next, a new formulation of partnership concurrency is proposed that both provides a natural stochastic generalization of deterministic social mixing functions, and allows for the analytic determination of equilibrium quantities of interest, such as the number of individuals with a given number of partners. Altmann [3,4] describes the distribution of the number of partners, given the population-level rate of adding new partners as a function of the current number of partners. That result is extended below by a derivation of population-level rates from individual-level parameters.

Section 3 generalizes the simplified model in three ways. First, partnership and infection durations can be modeled as arbitrary gamma distributed random variables, rather than requiring all quantities of interest to be exponentially distributed. Second, the parameters of individuals are allowed to vary from individual to individual. For instance, each individual may have a different propensity to partner or to acquire concurrent partners. Third, partnerships can be formed in a

variety of social or geographic settings, in a way that is motivated by the structured mixing formulation of Jacquez et al. [15]. The duration and sexual contact rate of partnerships may depend on the context in which the partnership is formed, so that, for instance, that long-term partnerships and partnerships with commercial sex workers can be modeled. This generalized process is expressible as a generalized semi-Markov process, and as such, is amenable for computer simulation. GERMS is a computer implementation of this general model.

Section 4 discusses discrete-event simulation experiments for the model. The results indicate that the large population approximations provide good approximations for populations of five hundred individuals. Graphs indicate how infection levels vary above and below pseudo-equilibrium levels. Simulations also quantify the change in endemic infection levels due to heterogeneity in individual parameters and partnership dynamics. The partnership-formation approximations are shown, for practical purposes, to reasonably represent the underlying mathematical model.

2. SIS model for simple homogeneous populations

Consider a population of N_M male and N_F female individuals, where a sexually transmitted disease is spread in the context of heterosexual partnerships. Individuals are presumed to be susceptible or infectious (SIS model), in the manner that gonorrhea has most commonly been modeled [26]. All individuals are monogamous at any given time, although individuals may form multiple partnerships sequentially. Transmission is presumed to occur only in the context of a partnership. Each individual is modeled as a node in a network. Recruitment and death adds and removes nodes from the network. For simplicity, recruitment and death are not included, so the population is closed. An arc indicates an ongoing partnership between individuals. An individual's infection status is a property of the corresponding node, as shown in Fig. 1. Arcs are added and removed stochastically through time as individuals form and dissolve partnerships.

As such, the population dynamics model is a network-valued stochastic process that is motivated by the deterministic model of Dietz and Hadeler [8]. The state space is a stochastic network

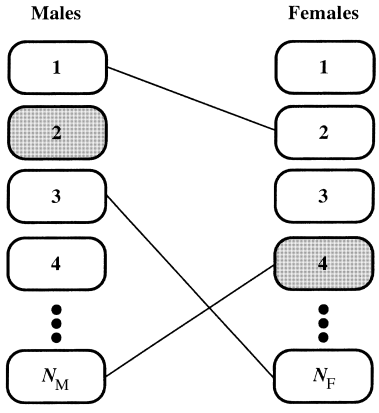

Fig. 1. A stochastic discrete-individual infection model that represents individuals with nodes and partnerships with arcs. Infected individuals are identified with shaded nodes.

(a set of nodes, an infection status for each node, and a set of arcs) whose state transition matrix depends both on properties of individuals, such as the partnership propensities of individuals that are unpartnered, and on parameters that depend on both partners, such as the duration of a partnership. The models of [11,20,24] are also discrete-event processes, and [20,24] take place in discrete-time. The continuous time (CT) process presented here can also be expressed as a discrete-time Markov chain with uniformization (e.g., see, [26]). However, our approach to retaining a continuous time formulation simplifies the derivation of analytic properties below, and provides a formal mechanism for studying the discrete-event simulation model. See also Kretzschmar [18] for a related approach to modeling.

An individual's propensity to form new partnerships is conceived as a concept related to, but distinct from, the rate at which two individuals actually form a partnership. Because of the monogamy condition, male j has a partnership propensity $\xi_{\mathrm{M}j}$ that equals 0 when he is partnered, and that equals the base partnership propensity $\lambda_{\mathrm{M}j}$ when he is not partnered. Let $\xi_{\mathrm{F}k}$ and $\lambda_{\mathrm{F}k}$ be analogous quantities for female k. Because partnership status changes through time, so too may the partnership propensities $\xi_{\mathrm{M}j}$ and $\xi_{\mathrm{F}k}$. Let g represent the gender in the subscript. Table 1 summarizes this notation and notation used below.

The instantaneous rate of partnership formation (formation of an arc) between j and k is determined by a mixing rate function r_{jk} that depends on the partnership propensities and the total numbers of individuals, $n_{\mathrm{M}j}$ and $n_{\mathrm{F}k}$, that are eligible to partner with male j and female k, respectively. One choice for the mixing rate function is the arithmetic mean of the partnership propensities after normalizing for the number of possible partners

$$r_{jk} = \begin{cases} \frac{1}{2}\left(\frac{\xi_{\mathrm{M}j}}{n_{\mathrm{M}j}} + \frac{\xi_{\mathrm{F}k}}{n_{\mathrm{F}k}}\right) & \text{if } j \text{ and } k \text{ are eligible to partner,} \\ 0 & \text{otherwise ,} \end{cases} \tag{1}$$

where 'eligible to partner' means that each individual currently has no partner. Appendix A.1 indicates that this mixing function is a special case of an individual-level analog of the Frederickson/McFarland [28,29] properties for population-level mixing functions. The sum of the r_{jk} over all potential couples (j, k) is the overall rate r of partnership formation

$$r = \sum_{j=1}^{N_{\mathrm{M}}}\sum_{k=1}^{N_{\mathrm{F}}} r_{jk} = \left(\sum_{j=1}^{N_{\mathrm{M}}} \frac{\xi_{\mathrm{M}j}}{2}\right) + \left(\sum_{k=1}^{N_{\mathrm{F}}} \frac{\xi_{\mathrm{F}k}}{2}\right) \tag{2}$$

The term $n_{\mathrm{M}j}$ drops out of Eq. (2) because $\xi_{\mathrm{M}j}$ is non-zero for exactly $n_{\mathrm{M}j}$ of the r_{jk} ($n_{\mathrm{F}k}$ similarly drops out). As $n_{\mathrm{M}j}$ and $n_{\mathrm{F}k}$ change through time, partnership formation is a non-homogenous Poisson process (e.g., see, [27]). This formulation does not make the chances that a male will form a partnership with different females proportional to the partnership propensity of those females (or vice versa). It has the nice quality, however, that it makes the overall rate r of heterosexual partnership formation in the population the average of the total partnership potential of each gender. If the partnership propensities are the same ($\lambda = \lambda_{\mathrm{M}j} = \lambda_{\mathrm{F}k}$) for all individuals in a sequentially monogamous population, then $n_{\mathrm{M}j}$ is the number of unpartnered males, $n_{\mathrm{F}k}$ is the number of unpartnered females, and $r = \lambda(n_{\mathrm{M}j} + n_{\mathrm{F}k})/2$.

Separation (removal of an arc) occurs randomly through time as well. The Markovian character of partnership dynamics is preserved by assuming that each partnership separates with rate σ, so that $1/\sigma$ is the mean duration of a partnership.

Table 1
Notation for network-valued stochastic process

Notation	
For defining the stochastic process	
$\lambda_{g,\ell}$	Base partnership propensity for an individual ℓ that has gender g (M or F)
$\xi_{g,\ell}$	Current partnership propensity for an individual ℓ that has gender g (M or F)
N_g	Total number of individuals that have gender g
$n_{g,\ell}$	Number of individuals that can partner with an individual ℓ that has gender g
σ	Separation rate for partnerships
h	Sex-act rate during partnerships
β	Transmission probability per exposure
d	Mean duration of infection (recovery rate $= 1/d$)
(x_a, y_a)	Number of (males, females) with infection status a (i for infected or u for uninfected)
z_{ab}	Number of couples where the male has infection status a and the female has infection status b
s	The vector $s = (x_i, y_i, x_u, y_u, z_{uu}, z_{iu}, z_{ui}, z_{ii})$
B	Number of social-geographic settings ('bins') for partnership formation
$f_{i,\ell}$	Fraction of partnership propensity of individual ℓ that is allocated to bin i
For approximating pseudo-equilibrium infection levels	
π_u	Fraction of unpartnered individuals that are infected
π_p	Fraction of partnered individuals that are infected.
$\tilde{\pi}_u, \tilde{\pi}_p$	Pseudo-equilibrium fraction of infected unpartnered and partnered individuals
$\tilde{\pi}$	Pseudo-equilibrium fraction of infected individuals in entire population
$g_{\bullet,\ell}$	Fraction of partnerships that terminate with exactly ℓ infected individuals
$g_{i,\ell}$	Fraction of partnerships that terminate with ℓ infected individuals, assuming that the partnership started with exactly i infected individuals
For determining the basic reproduction number	
R_0	Basic reproduction number
X	The expected number of new infections directly caused by an infected individual given that they have just ended a partnership
Y	The expected number of new infections directly caused by an infected individual given that they just started a partnership with an uninfected individual
Z	The expected number of new infections directly caused by an infected individual given that they just started a partnership with an infected individual
For evaluating partnership concurrency	
p_ℓ	Proportion of individuals with ℓ partners
$q_{\ell k}$	Number of partnerships between males with ℓ partners and females with k partners
α_ℓ	Rate that some individual in the population with ℓ partners forms an additional concurrent partnership

The infection process for an SIS model must account for two distinct features: transmission and recovery. Each individual (node) is either infected or uninfected. Transmission accounts for new infections, and the duration of infection is presumed to be exponential with mean d, so that the recovery rate is $1/d$. The rate may depend on the infected individual's gender. For instance, [26] uses $d_M = 10$ days for males and $d_F = 100$ days for females in a study of gonorrhea.

During each partnership, exposure events such as sex-acts are modeled as a Poisson process with rate h during the partnership, and the probability of transmission per exposure is β. Thus, if a partnership has duration τ, and it is known that one partner is infected and the other uninfected

when the partnership begins, and that the infected partner will not recover prior to the end of the partnership, the probability of transmission during the partnership is $1 - \exp[-h\beta\tau]$, the tail probability of an exponential distribution. This formulation presumes that a partnership is not necessarily initiated with a sex-act.

2.1. Sequentially monogamous populations

Suppose that there are the same number $N = N_M = N_F$ of males and females in a heterosexual, sequentially monogamous population. The discussion that follows can also be applied to disassortative mixing patterns in populations of individuals with distinguishing characteristics other than gender (e.g., distinct roles for homosexual partnerships). The state transition rate matrix for this model is quite large, since the state must represent all nodes and the presence and absence of each possible arc in the network. The entire state may be needed for evaluating the effectiveness of individual-based control measures, such as contact tracing [16].

However, the analysis can be simplified if only summary information is desired and the population is homogeneous. Suppose that each individual has the same base partnership propensity λ, all partnerships have the same sex-act rate h and separation rate σ, and that the recovery rate $1/d$ and probability of transmission per contact β are the same for both sexes. Let $s = (x_i, y_i, x_u, y_u, z_{uu}, z_{iu}, z_{ui}, z_{ii})$ be the state vector whose elements have the following meaning: x_i denotes the number of single and infected males, y_i the number of single and infected females, x_u the number of single and uninfected males, y_u the number of single and uninfected females, z_{uu} the number of couples with two uninfected partners, z_{iu} the number of couples where only the male partner is infected, z_{ui} the number of couples where only the female partner is infected, and z_{ii} denotes the number of couples where both partners are infected.

Given the simplifying assumptions of this section, this vector is also the state of a Markov chain whose transition rates are given in Table 2. The state s' is indicated by listing only those coordinates that are changed, and the table assumes that s is not a boundary state (i.e., one or more elements have 0 or N). For states s at the boundary, some of the transitions are not possible, and therefore have rate 0.

A few observations can be made. Let $z = z_{uu} + z_{iu} + z_{ui} + z_{ii}$ be the total number of partnerships, independent of infection status (so the range of z is 0 to N), and let $\pi_{z,\ell}$ be the probability that $z = \ell$. Given the assumptions made here in Section 2.1, the rate of change of $\pi_{z,\ell}$ can be determined from the first two lines of Table 2 to be

$$\begin{aligned}
\frac{d\pi_{z,0}}{dt} &= \sigma\pi_{z,1} - N\lambda\pi_{z,0},\\
\frac{d\pi_{z,N}}{dt} &= \lambda\pi_{z,N-1} - N\sigma\pi_{z,N},\\
\frac{d\pi_{z,\ell}}{dt} &= (N-(\ell-1))\lambda\pi_{z,\ell-1} + (\ell+1)\sigma\pi_{z,\ell+1} - (\ell\sigma + (N-\ell)\lambda)\pi_{z,\ell}
\end{aligned} \tag{3}$$

for $\ell = 1, \ldots, N-1$. The equilibrium distribution for the total number of partnerships z is found by setting $d\pi_{z,0}/dt = \cdots = d\pi_{z,N}/dt = 0$, and is verified by algebra to be binomial $(N, \lambda/(\lambda+\sigma))$. This implies that the distribution of the fraction z/N of individuals in a partnership has mean

Table 2
Transition rates from state $s = (x_i, y_i, x_u, y_u, z_{uu}, z_{iu}, z_{ui}, z_{ii})$ to s' (same as s except as noted) for an SIS infection in simplified population

s'	Rate: s to s'	Meaning
$x_a - 1, y_b - 1, z_{ab} + 1$	$\lambda(x_a + y_b)/2$	Partnership formed between male with infection status a and female with infection status b
$x_a + 1, y_b + 1, z_{ab} - 1$	σz_{ab}	Partnership separation between male with infection status a and female with infection status b
$z_{iu} - 1, z_{ii} + 1$	$h\beta z_{iu}$	Infection of female
$z_{ui} - 1, z_{ii} + 1$	$h\beta z_{ui}$	Infection of male
$x_u + 1, x_i - 1$	σx_i	Recovery of single male
$y_u + 1, y_i - 1$	σy_i	Recovery of single female
$z_{ub} + 1, z_{ib} - 1$	σz_{ib}	Recovery of male whose partner has infection status b
$z_{bu} + 1, z_{bi} - 1$	σz_{bi}	Recovery of female whose partner has infection status b
All other s'	0	Other state changes are not possible

$\lambda/(\lambda + \sigma)$ and variance $\lambda\sigma/[N(\lambda + \sigma)^2]$. The familiar deterministic limit as $N \to \infty$ has $\lambda/(\lambda + \sigma)$ of the population in a partnership.

Under the closed population assumption (even when the number of males and females are not the same), once nobody is infected, nobody will be infected ever again. Since the Markov chain is irreducible, the equilibrium probability π_s for state s is $\pi_s = 0$ for all states s such that someone is infected ($x_i + y_i + z_{iu} + z_{ui} + z_{ii} > 0$). Because of this, the equilibrium distribution is not relevant to studying infection levels for this model.

Pseudo-equilibrium infection levels, however, are non-zero infection levels for a stochastic infection model that bear a close relation to non-zero endemic levels for analogous deterministic differential equation infection models [3,4,30,31]. The pseudo-equilibrium distribution for the number infected in a population is the stationary distribution of the number infected, conditional on the event that absorption has not occurred. If the population size grows arbitrarily large, the normalized fraction of infectious individuals converges weakly to the non-zero endemic level predicted by the analogous differential equation model [30,31]. Here, pseudo-equilibrium levels are approximated for the above stochastic, finite population SIS model by using a large population approximation, as in [30]. The idea is to find states of the system so that the expected number of infected individuals in the population does not change through time, at least over an infinitesimally short time interval. See [3,4] or [31] for a more formal treatment of deterministic limits of related stochastic models.

To simplify the discussion, consider a homosexual (rather than heterosexual) population of $2N$ individuals that form sequentially monogamous partnerships. Each unpartnered individual forms partnerships with each other unpartnered individual at a rate equal to the sample average of the partnership propensities, as in Eq. (1). The sample prevalence levels presented here are also applicable to sequentially monogamous heterosexual populations with N individuals of each gender, assuming that both males and females have the same parameters, but the argument is more complex and does not provide additional insight.

Pseudo-equilibrium levels are determined here by identifying requisite constraints on the fraction π_u of *unpartnered* individuals that are infected and the fraction π_p of *partnered* individuals that are infected. Because the system is Markovian, π_u is also the fraction of individuals that are

infected at the moment they form a partnership. Similarly, π_p is the fraction of individuals that are infected at the moment they separate from a partnership.

First, the probability that someone infected at the end of a one partnership is still infected at the start of the next partnership is $\lambda/(\lambda + 1/d)$. This implies that a necessary condition for pseudo-equilibrium infection levels to be maintained is

$$\pi_u = \frac{\lambda}{(\lambda + 1/d)} \pi_p. \tag{4}$$

A second requirement is determined by examining the expected number of partners who are infected at the end of a partnership. Because the probability that each partner is infected at the beginning of a relationship is π_u, the number of infected individuals at the start of a new partnership is 0, 1, and 2 with probability $(1 - \pi_u)^2$, $2\pi_u(1 - \pi_u)$ and π_u^2, respectively (see Fig. 2).

If $g_{\bullet,\ell}$ is the fraction of partnerships that terminate with exactly ℓ infected individuals ($\ell = 0$, 1, 2), then the expected fraction of individuals that are infected when at the end of their partnerships is $(2g_{\bullet,2} + g_{\bullet,1})/2$. Appendix A.2 argues that this implies

$$\pi_p = \frac{2\pi_u(1 - \pi_u)\sigma h\beta + \pi_u^2\sigma(\sigma + h\beta + 1/d) + \pi_u(1 - \pi_u)\sigma(\sigma + 2/d) + \pi_u^2\sigma/d}{\sigma^2 + \sigma h\beta + 3\sigma/d + 2/d^2} \tag{5}$$

and further that these two relationships result in a quadratic equation for π_u with roots 0 (the equilibrium prevalence) and the pseudo-equilibrium endemic prevalence $\tilde{\pi}_u$ for unpartnered individuals,

$$\tilde{\pi}_u = \frac{\sigma(\sigma + 2h\beta + 2/d) - (\sigma^2 + \sigma h\beta + 3\sigma/d + 2/d^2)((\lambda + 1/d)/\lambda)}{\sigma h\beta}. \tag{6}$$

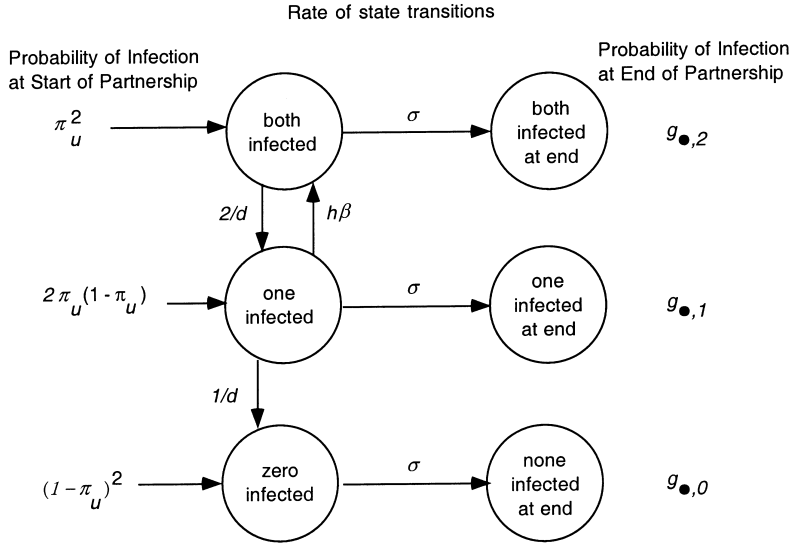

Fig. 2. State transition probabilities for the number of infected at the beginning and end of a partnership, used to approximate pseudo-equilibrium levels.

Infection is not sustained for an appreciable amount of time when $\tilde{\pi}_u$ in Eq. (6) is non-negative. The above argument is similar to the analysis by Dietz and Hadeler [8] that further includes birth and death processes, as well as more general mixing functions, but differs in that [8] assumes a relationship starts with a sexual contact and does not treat stochastic behavior.

From Eq. (4), the pseudo-equilibrium prevalence for partnered individuals is $\tilde{\pi}_p = \tilde{\pi}_u(\lambda + 1/d)/\lambda$. By taking a time-weighted average over the time spent in and out of partnerships, the pseudo-equilibrium fraction $\tilde{\pi}$ of infected individuals in the entire population is

$$\tilde{\pi} = \frac{\sigma}{\lambda+\sigma}\tilde{\pi}_u + \frac{\lambda}{\lambda+\sigma}\tilde{\pi}_p = \frac{\sigma+\lambda+1/d}{\lambda+\sigma}\tilde{\pi}_u. \tag{7}$$

Again, this result is based on the well-known results [3,30,31] for large population limits for pseudo-equilibrium distributions that show weak convergence to a deterministic equilibrium. These fractions should be interpreted as approximations for the mean pseudo-equilibrium prevalence, conditional on the event that infection has not been eliminated from the population.

The basic reproduction number, R_0, calculation for the above model must account for the possibility of the infection, recovery and reinfection of a partner during a single partnership. R_0 is evaluated by conditioning on the infection and partnership status of an individual. We use a conditional expectation argument that can be compared with the differential equation analysis of Dietz and Hadeler [8] for a similar infection model.

To this end, let X be the expected number of secondary infections caused by an unpartnered individual prior to recovery. Let Y be the expected number of secondary infections caused by an infected individual prior to recovery, given that the infected individual initially has an *uninfected* partner. Let Z be the expected number of secondary infections caused by an infected individual prior to recovery, given that the infected individual initially has an *infected* partner. Each of X, Y, Z is a function of the fraction π_u of infected individuals in the unpartnered population, as well as the infection system parameters.

Fig. 3 presents a tree that shows what events can happen from the time of infection of the individual until the end of the first partnership entered after infection. The transition probabilities are also shown. The reproduction number can be derived by examining how individuals change states through the tree in Fig. 3. The reproduction number is defined here as the number of

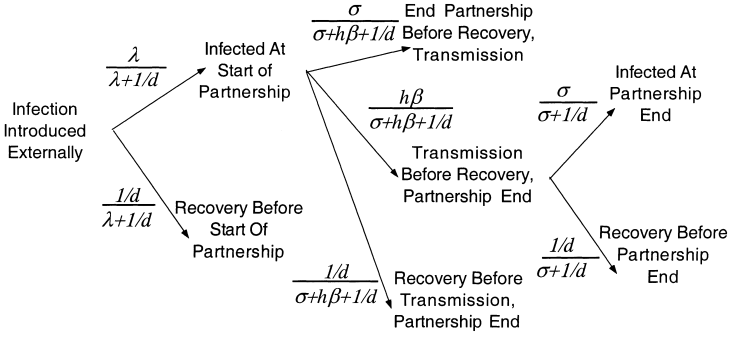

Fig. 3. State transition probabilities for simplified SIS model after introducing infection to an unpartnered individual, used to evaluate R_0.

secondary transmissions caused by an initial individual before recovery. The SIS assumption, plus the fact that partnerships have duration, implies that the initially infected individual may infect a partner who then recovers and becomes reinfected before the initially infected individual recovers. In our model, the infection, recovery, and reinfection of a single partner counts as two secondary transmissions, assuming that the initially infected individual had not recovered before the partner became reinfected.

There are several possible state changes for an infected, unpartnered individual (the index case). He/she recovers prior to forming a partnership with probability $1/d/(\lambda + 1/d)$. With probability $\lambda/(\lambda + 1/d)$, a partnership is formed first. The partner is infected with probability π_u (leading to Z additional secondary infections in expectation prior to the recovery of the index case) and uninfected with probability $1 - \pi_u$ (leading to Y additional secondary infections in expectation). This leads to the following relationship for X in terms of Y and Z.

$$X = \frac{1/d}{\lambda + 1/d}0 + \frac{\lambda}{\lambda + 1/d}[\pi_u Z + (1 - \pi_u)Y]. \tag{8}$$

The three possibilities for an infected individual whose partner is infected are (a) the individual recovers, (b) the partner recovers (and is therefore susceptible to reinfection), or (c) the partnership ends. These three outcomes lead to three terms in the following relationship for Z:

$$Z = \frac{1/d}{\sigma + 2/d}0 + \frac{1/d}{\sigma + 2/d}Y + \frac{\sigma}{\sigma + 2/d}X. \tag{9}$$

Similarly, an individual will infect an uninfected partner with probability $h\beta/(\sigma + h\beta + 1/d)$. This generates one new secondary case and has the potential of generating Z additional secondary infections, in expectation. Separation occurs prior to separation and recovery with probability $\sigma/(\sigma + h\beta + 1/d)$, and generates X more secondary infections in expectation. This implies that

$$Y = \frac{\sigma}{\sigma + h\beta + 1/d}X + \frac{h\beta}{\sigma + h\beta + 1/d}[1 + Z] + \frac{1/d}{\sigma + h\beta + 1/d}0. \tag{10}$$

Appendix A.3 shows that

$$X = \frac{d\lambda h\beta}{(\lambda + \sigma + 1/d)(\sigma + h\beta + 2/d)}[(\sigma + 1/d)(1 - \pi_u) + 1/d]. \tag{11}$$

If $h\beta$ approaches infinity (transmission is certain) and $\pi_u = 0$ (e.g., at the beginning of an epidemic) then X becomes $\lambda(d\sigma + 2)/(\lambda + \sigma + 1/d)$. If it were not possible to reinfect a recovered partner during one long partnership, then a similar analysis indicates that X would be $\lambda(d\sigma + 1)/(\lambda + \sigma + 1/d)$. The expected number of secondary infections therefore exceeds the expected number of individuals that become infected, because there is a chance that some individuals will be infected and reinfected by the same partner.

Even when limits for $h\beta$ and σ are not taken, the pseudo-equilibrium prevalence $\tilde{\pi}_u$ in the unpartnered population satisfies $Z(\tilde{\pi}_u) = 1$, meaning that the expected number of new infections generated by each infection before recovery is one. Therefore $Z(\pi_u)$ plays the role of the reproduction number R of the infection. The basic reproduction number R_0 is somewhat more complex in the context of stochastic processes than in the deterministic model context where contacts have no duration. The reason is that one must specify what one really means by an individual starting

out in an otherwise completely susceptible population of a *large* number of individuals. If the infected individual is unpartnered, say, by becoming infected by an individual that is external to the modeled population, then $R_0 = X(0)$, where $X(\pi_u)$ indicates a dependence on the prevalence π_u in the unpartnered population. If the individual is partnered with an uninfected individual, then $R_0 = Y(0)$. A reasonable assumption [3] is to define $R_0 = Z(0)$ as the expected number of secondary cases, given that a newly infected individual was infected by a partner, but that the rest of the population is susceptible. This definition is used in the remainder of the paper.

Note that $\tilde{\pi}$ does *not* equal any of $1 - 1/X(0)$, $1 - 1/Y(0)$, $1 - 1/Z(0)$. Since contacts occur at rate h during the fraction of time $\lambda/(\lambda + \sigma)$ that an individual is in a relationship, the effective contact rate is $c = h\lambda/(\lambda + \sigma)$, and $c\beta d = [h\lambda/(\lambda + \sigma)]\beta d$. Further note that $\tilde{\pi}$ does not equal $1 - 1/c\beta d$ (the well-known formula for prevalence that is derived by assuming that contacts are part of a point process, and that reinfection is not applicable). Empirical results in Section 4.1 indicate that $\tilde{\pi}$ is between $1 - 1/Z(0)$ and $1 - 1/c\beta d$. The reason for the discrepancy is that contacts are no longer a sequence of independent events, but are correlated as a result of the partnership dynamics. This discrepancy is consistent with the general conclusions of Kretzschmar and Dietz [19] in a study the role of partnership duration for a deterministic infection model, but the specific numbers differ as they use a different mixing function. At a general level, their model [19] also indicates that the prevalence may decrease under certain situations even when $R(0) > 1$. The same may happen for the stochastic SIS model above, in that the prevalence may initially decrease, in expectation, if all infectious individuals are initially unpartnered.

2.2. Populations with partnership concurrency

Stochastic models of concurrency are perceived as highly relevant for the analysis of infection levels in a population. Kretzschmar and Morris [20] link the distribution of the number of partners in a large population of individuals to epidemiological quantities of interest. They also present a discrete-time discrete-individual model of infection, and empirically compare network statistics to prevalence, but do not derive the equilibrium distribution of the number of partners as a function of the parameters of their discrete-individual model. Altmann [3] describes how to determine R_0 for an arbitrary distribution of the number of partners in a homogeneously mixing population with an SIR infection, and also derives the equilibrium distribution of the number of individuals with a given number of partners. Whereas [3] assumes that the rate of forming an additional partnership (as a function of the current number of partners) is a known, population-level parameter, its value is derived below from individual-level parameters. This section also derives R_0 in the sense of [3] for the SIS model with the potential for infection, recovery and reinfection during a *single* partnership.

As in Section 2.1, let λ, σ, h, β be the same for all individuals and partnerships, and let $N = N_M = N_F$. Let the partnership propensity $\xi_{g,\ell}$ of individual ℓ of gender g be damped by a factor $\theta \in (0, 1)$ as an individual increases his or her number of partners, $m_{g,\ell}$, so that

$$\xi_{g,\ell} = \lambda\theta^{m_{g,\ell}}. \tag{12}$$

The factor θ represents an individual's preference for seeking additional concurrent partners. If 0^0 is defined to be 1, then the sequential monogamy of previous sections can be represented by $\theta = 0$. The rate that a partnership is formed between male j and female k can still be defined as in Eq. (1),

and the overall rate of partnership formation is given in Eq. (2). The number of individuals available for partnering is $n_{g,\ell} = N_{\bar{g}} - m_{g,\ell}$ when everyone has a common $\theta \in (0,1)$, and $\bar{g}$ is the opposite gender, so that $n_{g,\ell}$ is approximately $N_{\bar{g}}$ when concurrency is allowed for everyone, as opposed to approximately $N_{\bar{g}}\sigma/(\lambda+\sigma)$ when all are sequentially monogamous. (Technically, summands in Eq. (2) corresponding to the $\xi_{g,\ell}$ of individuals that are already in partnerships with *everyone* of the opposite gender should be excluded.)

Suppose that p_ℓ is the fraction of individuals in a large population with exactly ℓ partners. The same p_ℓ is used for both genders because of the symmetry in partnering and infection parameters, and to simplify the arguments. The equilibrium value of p_ℓ can be approximated by determining large population limits for the average partnership propensity and break-up rate, as a function of the number of concurrent partners of an individual. From Eqs. (1) and (12), the rate at which a specific male with ℓ partners forms a partnership with a specific female with k partners is

$$r_{\ell k} = \frac{1}{2}\left(\frac{\lambda\theta^\ell}{N-\ell} + \frac{\lambda\theta^k}{N-k}\right).$$

By assumption, Np_ℓ males have ℓ partners, and Np_k females have k partners. Let $q_{\ell k}$ be the number of partnerships between males with ℓ partners and females with k partners. There are $Np_\ell Np_k - q_{\ell k}$ possible partnerships that can form between a male with ℓ partners and a female with k partners. The rate that such a partnership is formed is therefore

$$\frac{Np_\ell Np_k - q_{\ell k}}{2}\left(\frac{\lambda\theta^\ell}{N-\ell} + \frac{\lambda\theta^k}{N-k}\right).$$

Sum over k to get the rate α_ℓ that a male with ℓ partners forms a partnership with anyone,

$$\alpha_\ell = \sum_{k=0}^{N-1} \frac{(Np_\ell Np_k - q_{\ell k})(\lambda\theta^\ell/(N-l) + \lambda\theta^k/(N-k))}{2}.$$

It is useful to define the following large population approximation $\tilde{\alpha}_\ell$ for α_ℓ:

$$\tilde{\alpha}_\ell = Np_\ell \frac{\lambda\theta^\ell + \gamma}{2} \approx \alpha_\ell, \tag{13}$$

where $\gamma = \sum_{i=0}^{N-1} \lambda\theta^i p_i$ is a large population approximation for the average partnership propensity. As N increases without bound (scale $q_{\ell k}$ proportional to N to maintain a constant proportion of partnerships between males with ℓ partners and females with k partners) the following are true. The infinite sum γ converges to a finite value because $\theta < 1$. Both α_ℓ/N and $\tilde{\alpha}_\ell/N$ converge to $p_\ell(\lambda\theta^\ell + \gamma)/2$. And the total partnership formation rate, normalized by the population size, converges to the large population limit $\sum_{\ell=0}^{N-1} \alpha_\ell/N \to \gamma$.

The equilibrium fraction p_ℓ of individuals with ℓ partners can then be determined. In equilibrium, the overall break-up rate $Np_\ell(\ell\sigma)$ for the Np_ℓ males with ℓ partners must equal the overall partnership formation rate $\alpha_{\ell-1}$ for the $Np_{\ell-1}$ males with $\ell-1$ partners. The large population approximation $\tilde{\alpha}_\ell$ for α_ℓ implies that

$$\frac{p_\ell}{p_{\ell-1}} = \frac{\lambda\theta^{\ell-1} + \gamma}{2\ell\sigma} \quad \text{for } \ell = 1, \ldots \tag{14}$$

The additional constraint $\sum p_\ell = 1$ then determines each p_ℓ. The p_ℓ can then be used to determine epidemiological quantities of interest, such as the rate of epidemic rise.

If R_0 is interpreted as the expected number of secondary cases produced by a 'typical' case early in the epidemic, an argument similar to that found in [3] can be modified to determine R_0 in the context here. Define η_ℓ to be the average number of secondary cases produced by an initial case with ℓ partners. If one samples individuals with probability that is weighted by the number of partners, then an individual with ℓ partnerships is sampled randomly with probability $\ell p_\ell / \sum_j j p_j$, and

$$R_0 = \frac{\sum_\ell \ell p_\ell \eta_\ell}{\sum_\ell \ell p_\ell}.$$

Assume that the initial case in question is infected by one infected partner, and that the other $\ell - 1$ partners are initially uninfected. Let χ be the expected number of times that each of the $\ell - 1$ susceptible partners will become infected (including reinfections) before the initial case recovers. Those susceptible partners will be infected with probability $h\beta/(h\beta + \sigma + 1/d)$. Further, the probability that an infected partner recovers before the end of a partnership is $(1/d)/(2/d + \sigma)$. Because of the Markovian assumptions, a partner that recovers is reinfected by the initial case another χ times, in expectation, if the initial case is still infected, and 0 times, if the initial case also recovered. These observations imply that

$$\chi = \frac{h\beta}{h\beta + \sigma + 1/d}\left(1 + \frac{1/d}{2/d + \sigma}\chi\right).$$

so that

$$\chi = \left(\frac{h\beta + \sigma + 1/d}{h\beta} - \frac{1/d}{2/d + \sigma}\right)^{-1}.$$

Take an expectation over the number of susceptible partners at the time of transmission to obtain the expected number of infection transmissions to susceptible individuals that were partnered with the base case at the time of infection

$$\chi \frac{\sum_\ell (\ell - 1)\ell p_\ell}{\sum_\ell \ell p_\ell}.$$

A recurrence relation for the number of additional partnerships v_ℓ formed before recovery by a person with ℓ partners can be derived by conditioning on whether recovery, partnership formation, or separation occurs first:

$$v_0 = \frac{1/d}{\alpha_0 + 1/d}0 + \frac{\alpha_0}{\alpha_0 + 1/d}(1 + v_2),$$

$$v_\ell = \frac{1/d}{\alpha_\ell + 1/d + \ell\sigma}0 + \frac{\alpha_\ell}{\alpha_\ell + 1/d + \ell\sigma}(1 + v_{\ell+1}) + \frac{\ell\sigma}{\alpha_\ell + 1/d + \ell\sigma}v_{\ell-1}, \quad \text{for } \ell \geqslant 1.$$

Each new partner becomes infected χ times in expectation. If new partners are presumed to be uninfected, the expected number of new infections to new partners is $\chi \sum_\ell \ell p_\ell v_\ell / \sum_\ell \ell p_\ell$. The basic reproduction number is therefore

$$R_0 = \chi \frac{\sum_\ell \ell p_\ell [v_\ell + (\ell - 1)]}{\sum_\ell \ell p_\ell}. \tag{15}$$

Again, the interpretation is the expected number of new infections during the initial stages of infection. This derivation does not conclude that endemic prevalence is $1 - 1/R_0$. Equations for the pseudo-equilibrium prevalence for this concurrency model have not yet been derived.

3. A more general infection transmission system

The simple model in Section 2 ignores a number of real-world complexities that may affect the prevalence of infection. This section describes three model extensions that relax initial assumptions in order to improve the model's ability to approximate reality. Each extension has been implemented in the GERMS simulation model described in [2,16].

Infection and partnership durations might not be represented best by exponential distributions. Infection durations are therefore allowed to have a gamma (d_α, d_β) distribution, so that the mean d_α/d_β and variance d_α/d_β^2 can be better controlled to suit modeler preference. Partnerships are also allowed to have a duration that has a gamma $(\sigma_\alpha, \sigma_\beta)$ distribution.

Heterogeneity of individuals may be needed to model individuals with different propensities for partnering and concurrency. All individual-related parameters are therefore allowed to vary from person to person. For instance, the base partnership propensities $\lambda_{\mathrm{M}j}$, $\lambda_{\mathrm{F}k}$ and concurrency attenuation factors $\theta_{\mathrm{M}j}$, $\theta_{\mathrm{F}k}$ are allowed to vary from person to person. Sequentially monogamous individuals have $\theta = 0$.

Heterogeneous mixing is required to account for the influence of social and geographic mixing patterns on infection dynamics. Jacquez et al. [15] proposed a structured mixing model that describes the rate at which members of different social groups meet in different activity settings. This motivates the inclusion of activity settings (called a 'bins' in [2,16]) in the present model. Each activity setting b (for $b = 1, \ldots, B$) has a geographical location (x, y coordinate) and social contexts. Social contexts are represented by allowing only certain subsets of individuals to form partnerships in each activity setting. Example contrasts for social contexts include core/non-core, heterosexual/homosexual, and well-insured/poorly-insured individuals. Each activity setting may be the site of partnership formation for individuals from one or more social groups, and each individual may be a member of one or more social groups.

To implement this form of heterogeneous mixing, a parameter $f_{b,\mathrm{M}j}$ is used to describe the fraction of the partnership propensity for male j allocated to activity setting b. Male j is only allowed to form partnerships in activity setting b if $f_{b,\mathrm{M}j} > 0$. The partnership propensity $\xi_{b,\mathrm{M}j}$ of male j in activity setting b is defined to be

$$\xi_{b,\mathrm{M}j} = f_{b,\mathrm{M}j}\lambda_{\mathrm{M}j}\theta_j^{n_{\mathrm{M}jj}}.$$

Similar notation, $\xi_{b,\mathrm{F}k}, f_{b,\mathrm{F}k}$, is used for female k. The choice of $f_{b,\mathrm{M}j}$ and $f_{b,\mathrm{F}k}$ can be guided by the geographic closeness of individuals to activity settings, and may depend on an individual's social context.

Partnering continues in each activity setting along the lines of the model presented in Section 2.1, except that the additional subscript b must be added. In particular, the instantaneous rate that male j and female k form a partnership in activity setting b is

$$r_{b,jk} = \begin{cases} \frac{1}{2}\left(\frac{\xi_{b,\mathrm{M}j}}{n_{b,\mathrm{M}j}} + \frac{\xi_{b,\mathrm{F}k}}{n_{b,\mathrm{F}k}}\right) & \text{if } j \text{ and } k \text{ are eligible to partner in } b, \\ 0 & \text{otherwise.} \end{cases} \tag{16}$$

Here, 'eligible to partner in b' means that (i) j and k are not already partnered to each other, (ii) both j and k can form partnerships in activity setting b, and (iii) both j and k either allow for partnership concurrency or are sequentially monogamous but unpartnered. Define the indicator variables $c_{b,jk}$ and $e_{b,\mathrm{M}j}$ by

$$c_{b,jk} = \begin{cases} 1 & \text{if } j \text{ and } k \text{ are eligible to partner in activity setting } b, \\ 0 & \text{otherwise,} \end{cases}$$

$$e_{b,\mathrm{M}j} = \begin{cases} 1 & \text{if there is at least one female eligible to partner with male } j \text{ in activity setting } b, \\ 0 & \text{otherwise.} \end{cases}$$

The rate $r_{b,j}$ that male j forms a partnership in activity setting b is therefore

$$r_{b,j} = \sum_{k=1}^{N_\mathrm{F}} r_{b,jk} = \frac{e_{b,\mathrm{M}j}}{2}\left(\xi_{b,\mathrm{M}j} + \sum_{k=1}^{N_\mathrm{F}} \frac{c_{b,jk}\,\xi_{b,\mathrm{F}k}}{n_{b,\mathrm{F}k}}\right). \tag{17}$$

The factor $e_{b,\mathrm{M}j}$ ensures that an individual with no eligible partners will not form a partnership, and $n_{b,\mathrm{M}j}$ drops out of the denominator in the first term on the right-hand side of Eq. (17) because there are $n_{b,\mathrm{M}j}$ values of k such that $c_{b,jk} = 1$. The probability that individual j is selected to participate in a partnership formed in activity setting b is proportional to $r_{b,j}$. The rate r_b of partnership formation in activity setting b is

$$r_b = \sum_{j=1}^{N_\mathrm{M}}\sum_{k=1}^{N_\mathrm{F}} r_{b,jk} = \left(\sum_{j=1}^{N_\mathrm{M}} \frac{e_{b,\mathrm{M}j}\,\xi_{b,\mathrm{M}j}}{2}\right) + \left(\sum_{k=1}^{N_\mathrm{F}} \frac{e_{b,\mathrm{F}k}\,\xi_{b,\mathrm{F}k}}{2}\right). \tag{18}$$

The overall partnership formation rate is therefore $r = r_1 + \cdots + r_B$. To complete the specification of the partnership dynamics for this general model, the duration of partnerships formed in activity setting b is assumed to have a γ distribution whose parameters $(\sigma_{b,\alpha}, \sigma_{b,\beta})$. This allows for long-term partnerships to be formed in one activity setting, and commercial sex partnerships to be formed in another. For short-term relationships in activity setting b, the sex-act rate parameter h_b may need adjustment to ensure that the probability of transmission per short-term partnership is modeled correctly, as the GERMS model does not presume that a relationship begins with a sex act.

4. Simulation experiments

Simulation experiments were employed to empirically observe the stochastic variation of the infection process around pseudo-equilibrium levels, and to check if the large-population limits are reasonable for small population sizes. See [2] for a further discussion of the GERMS simulation environment used to implement these experiments, including a description of an approximation for the partnership formation process that improves the runtimes while preserving the appropriate asymptotic limits.

4.1. Sequentially monogamous populations

The first set of simulation experiments examines relationships between R_0, the theoretical pseudo-equilibrium prevalence, and simulated endemic values for the monogamous population with homogeneous personal and mixing characteristics. The population is closed, with $N = 1000$ males and $N = 1000$ females that seek sequentially monogamous partnerships in a single activity setting. The 'base model' assumes that all individuals have a common partnership propensity $\lambda = 1/14$ days, separation rate $\sigma = 1/14$ days, contact rate during partnerships $h = 3/7$ (3 per week), with per-contact transmission probability $\beta = 0.3$ when exactly one partner is infected. The infection duration is exponentially distributed with mean $d = 55$ days, a value that is reasonable for untreated gonorrhea. Pseudo-equilibrium values were observed by (i) simulating partnership dynamics without infection for 1 year to attain close-to-equilibrium partnership dynamics, (ii) infecting approximately 20% of the population, (iii) simulating five more years allow transient effects to dissipate, then (iv) collecting prevalence statistics over an additional 15 years of simulated time. Similar experiments varied a single parameter from the base case.

The time-average of the prevalence during that 15 years is presented in Table 3, along with a batch-mean [32] confidence interval (30 batches of 6 months each). Statistical tests of independence for the batch means indicate that autocorrelation is insignificant, except for a small positive one-step correlation for the base case. For the base case, then, the CI may be slightly overconfident. The simulated batch mean compares quite well with the theoretical pseudo-equilibrium prevalence of Eq. (7). Table 3 also presents the expected number of secondary infections ($R_0 = Z(0)$) in the terminology used in the end of Section 2.1) due to an individual that was infected by a partner at the beginning of the epidemic, as well as $1 - 1/R_0$. Clearly the endemic, pseudo-equilibrium level does not equal $1 - 1/R_0$, using that definition of R_0. The endemic level is under-predicted by that equation. Nor does the pseudo-equilibrium level equal $1 - 1/(c\beta d)$, the endemic level predicted by deterministic models which have exposures occurring as point contacts with no duration. The actual endemic level is rather lower than predicted by $1 - 1/(c\beta d)$. The discrepancy is due to the fact that contacts are not independent events at specific points in time, but in fact are part of a correlated sequence of events that take place in the context of sexual partnerships, a sequence that further depends on the infection status of each partner.

The lack of correspondence between the expected number of secondary infections R_0, an individual-level concept, and population levels of infection is similar to what we have noted in other

Table 3
Comparison of basic reproduction number, the prevalence predicted by a theoretical model, and the simulated prevalence for an SIS infection in a homogeneous, monogamous population

Experiment	$R_0 = Z(0)$	$1 - 1/R_0$	$1 - 1/(c\beta d)$	Pseudo-equilibrium prevalence, $\tilde{\pi}$	Observed mean prevalence (simulation)	95% confidence intervals (simulation)
Base case	1.25	0.200	0.717	0.298	0.298	(0.288, 0.308)
Base, but $\lambda = 1/7$	1.69	0.408	0.788	0.570	0.570	(0.563, 0.577)
Base, but $\sigma = 1/21$	1.22	0.180	0.764	0.307	0.303	(0.291, 0.315)
Base, but $\beta = 0.5$	1.53	0.346	0.830	0.515	0.513	(0.507, 0.518)
Base, but $d = 50$	1.14	0.123	0.689	0.190	0.199	(0.189, 0.209)

non-random mixing simulations [17]. A point time proportionate mixing process is needed for the correspondence to be valid.

4.2. Populations with partnership concurrency

A second set of simulation experiments examines how the prevalence is affected by changing some of the population and infection dynamics to take advantage of the generality described in Section 3. Table 4 summarizes how the base case population was modified for each simulated population. The average prevalence is calculated over a 15-year period, as for the first set of experiments.

Table 4
Summary of second set of simulations (population of 2000 individuals) to evaluate prevalence as model parameters are changed (Fig. 4 displays the simulated prevalence)

Name of run	Description	Average prevalence
Base case	1000 males, 1000 females; Common $\lambda = 1/14$; common $\sigma = 1/14$; recovery rate $1/d = 1/55$; sequential monogamy ($\theta = 0$); sex-act rate $h = 3/7$; transmission probability per act $\beta = 0.3$.	595
Concurrency for all	Same as 'base case', except that everyone can have concurrent partnerships ($\theta = 0.1$)	993
Concurrency for 10%	Same as 'base case', except that exactly 10% of the males and females can have concurrent partnerships ($\theta = 0.1$)	710
Extra activity setting	Same as 'concurrency for 10%', except that monogamous individuals form partnerships exclusively in activity setting 1, and others split the partnership propensity in two settings	725

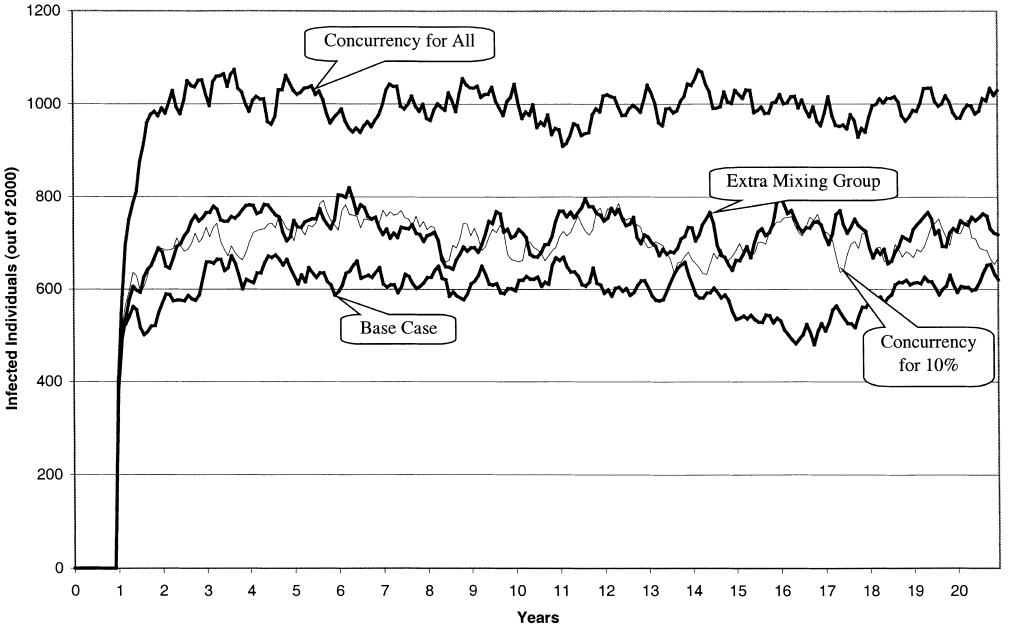

Fig. 4. Prevalence for several simulated scenarios (Table 4 describes the scenarios further).

Fig. 4 displays the prevalence, as a function of time, for each scenario. The 'base case' for this second set of experiments is the same as the base case for the first set of experiments. The prevalence clearly varies through time around the pseudo-equilibrium, after infection was introduced at time 12 months in the simulation. Partnership concurrency plays a significant role. If everyone may have concurrent partnerships ('concurrency for all'), with a relatively low level of concurrency ($\theta = 0.1$), then the prevalence increases to 993 from 595. Even if only 10% of the population have the ability to form concurrent partnerships ('concurrency for 10%'), the prevalence is raised significantly over the base case (710 vs 595). Prevalence was also examined under the assumption that individuals that form concurrent partnerships tend to form partnerships with each other, rather than mixing homogeneously with the entire population ('extra mixing group'). The mixing was moderately non-homogeneous (all monogamous individuals form partnerships exclusively in activity setting 1, and the 10% of the population that can have concurrent partnerships split their partnership propensity equally between activity setting 1, and a second activity setting reserved exclusively for individuals that can form concurrent partnerships). This non-homogeneity changed the prevalence very mildly compared to the importance of concurrency alone.

4.3. Evaluation of concurrency approximations

The last simulation experiment evaluates the approximations in Section 2.2 that predict the expected number of individuals in a population with a given number of partners. Iterative methods were used to compute the solutions for the relevant equations from Section 2.2. Both numerical error and the fact that a large-population approximation is used introduce the potential for inaccuracies into prediction.

The simulated population consisted of $N = 1000$ males and $N = 1000$ females with a common partnership propensity $\lambda = 1/14$ days, separation rate $\sigma = 1/14$ days, contact rate during partnerships $h = 3/7$ (3 per week), with per-contact transmission probability $\beta = 0.3$, and a concurrency damping factor $\theta = 0.3$. The recovery rate was $1/d = 1/55$. Both the theoretical analysis of Section 2.2 and the simulated estimates of the number of individuals with a given number of partners are presented in Table 5. Individuals with five or more partners are grouped together due to the small numbers associated with that level of concurrency. The approximations do quite well, although the theoretical values typically lie a very small amount outside the 95% confidence intervals for each quantity. This minor deviation is probably best explained as a result of the ap-

Table 5
Comparison of simulation estimates with the expected number of individuals with a given number of partners in equilibrium (out of a population of 2000)

Number of concurrent partners	Theoretical number of individuals (approximation)	Simulated number of individuals (95% CI from batch means)
0	988.5	(983.3, 988.5)
1	802.5	(801.8, 802.2)
2	185.2	(185.3, 185.8)
3	22.0	(22.3, 22.4)
4	1.8	(1.74, 1.76)
5 or more	0.11	(0.119,0.131)

proximations for the partnership generation rate described in [2], as the other large-population approximations all seem to hold extremely well for the other quantities of interest for sequentially monogamous populations.

For this experiment with partnership concurrency, R_0 in Eq. (15) evaluates to 2.223. The pseudo-equilibrium number of infected individuals estimated by the simulation is $1136.1 + / - 8.78$ out of 2000. This quantity is near, but several standard deviations higher than, $2000(1 - 1/2.223) \approx 1100$. If individuals could not be infected twice by the same partner during a given relationship, but were otherwise susceptible to reinfection by other partners during other infections, then an analysis similar to that in Section 2.2 indicates that R_0 would be 2.004. This corresponds to a prevalence of 1002 out of 2000. This illustrates the potential for reinfection during a partnership to influence the prevalence of infection.

5. Discussion and conclusions

The stochastic, discrete-individual infection transmission model presented above has parameters that correspond to deterministic model parameters for homogeneous populations. Because of this, a number of analytical results are available to describe equilibrium partnership dynamics and endemic infection levels. Moreover, the simulation implementation of the generalized version of the model allows one to study the effects of relaxing the homogeneity assumptions. This suggests that model exploration can proceed in stages, with deterministic models being used to explore some high-level issues. The general stochastic discrete-individual model can then be used to refine or extend the analysis. Potential uses include the study of stochastic variation in endemic infection levels through time, the evaluation of heterogeneity effects in the partnership formation process, and the observation of the effects of a non-exponential infection duration on the prevalence of infection. A further use is the exploration of individual-level social contact patterns and their role in the transmission of infection [16].

The particular heterosexual mixing formulation presented above has some desirable properties but it is not as general as the formulations presented by Castillo-Chavez and co-authors [5,6]. A major advantage of the specific formulation here is that analytical equilibrium results are available to help validate the simulation implementation of the model. The paper also derives population-level parameters for partnership concurrency from individual-level parameters, which may improve the ability of a modeler to set model parameters in a range consistent with field observations.

Contact processes that have partnerships with any duration cause a dissociation between the individual-level definition of R_0 and population-level consequences that are associated with R_0 in models with point-in-time random mixing processes. It is therefore inappropriate to use the basic reproduction number to determine both the rate of epidemic rise and the endemic level of infection when point-time contacts are not a valid model, even when a number of standard homogeneous population assumptions are valid. The reason is that infection opportunities are not independent events when contacts occur in the context of partnerships that have a non-zero duration. The potential for infection, recovery and reinfection of an individual that is in regular contact with an infected partner is also shown to increase the endemic prevalence of infection in a population.

Discrete-individual simulation experiments in continuous time indicate that the large-scale approximations are useful for populations with just a few thousand individuals, in the sense that

batch mean confidence intervals for the endemic levels contain the theoretical endemic prevalence. Additional experiments (results not shown) indicate that the approximations are also reasonable with 500 individuals.

These results help to bridge the gap between deterministic differential equation analysis and stochastic discrete-individual simulation models for the study of infectious diseases. There remains much more to be done, but the above mathematical formulation and insights for computer implementation may serve as a useful basis for further research.

Acknowledgements

The authors are pleased to acknowledge the financial support of the Centers for Disease Control (CDC), Atlanta. Some of this material was presented at the CDC meeting on STD Models in Santa Fe, NM on April 28, 1999. The assistance of Szu Hui Ng and Peter Yu is greatly appreciated. The authors gratefully acknowledge the computing resources provided by the Center for the Study of Complex Systems (CSCS) at the University of Michigan in the completion of this research. The constructive comments of two anonymous referees are also appreciated.

Appendix A

A.1. Mixing functions

A significant amount of work has been done to develop mixing rate functions for deterministic models [5,6,8]. The GERMS model formulation builds upon that work by applying the mixing rate to the individual, rather than population level. Consider, for example, the Fredrickson–McFarland properties [28,29] for a deterministic mixing rate $\varphi(n_{\mathrm{M}j},\ n_{\mathrm{F}k})$ when there are $n_{\mathrm{M}j}(n_{\mathrm{F}k})$ males (females) available for partnering:

(i) $\varphi(0, n_{\mathrm{F}k}) = \varphi(n_{\mathrm{M}j}, 0) = 0$ (no partnership formation if no males or females are available),

(ii) $\alpha\varphi(n_{\mathrm{M}j}, n_{\mathrm{F}k}) = \varphi(\alpha n_{\mathrm{M}j}, \alpha n_{\mathrm{F}k})$ (rate of formation is proportional to the population size),

(iii) $\varphi(n_{\mathrm{M}j} + a, n_{\mathrm{F}k} + b) > \varphi(n_{\mathrm{M}j}, n_{\mathrm{F}k})$ when $a, b > 0$. (rate is non-decreasing in population size).

GERMS also has the following individual-level characteristic: if the base partnership propensity of all individuals is multiplied by α, then the rate of partnership formation is multiplied by α, too.

By reinterpreting the population-level concept φ with the pair-level concept r_{jk}, and letting the number of individuals $n_{\mathrm{M}j}$ and $n_{\mathrm{F}k}$ be replaced by the normalized partnership propensities $\xi_{\mathrm{M}j}/n_{\mathrm{M}j}$ and $\xi_{\mathrm{F}k}/n_{\mathrm{F}k}$, respectively, one obtains alternate mixing rates analogous to deterministic population-level mixing functions already used in demographic studies [6]. Assuming that j and k can form a partnership, the arithmetic average of Eq. (1) might be substituted with $r_{jk} = \min(\xi_{\mathrm{M}j}/n_{\mathrm{F}}, \xi_{\mathrm{F}k}/n_{\mathrm{M}})$; the harmonic mean

$$r_{jk} = (\xi_{\mathrm{M}j}\xi_{\mathrm{F}k}/n_{\mathrm{F}}n_{\mathrm{M}})/(\xi_{\mathrm{M}j}/n_{\mathrm{F}} + \xi_{\mathrm{F}k}/n_{\mathrm{M}}),$$

or the geometric mean $r_{jk} = (\xi_{\mathrm{M}j}\xi_{\mathrm{F}k}/n_{\mathrm{F}}n_{\mathrm{M}})^{1/2}$. A mixing rate determined by the partnership propensity of females alone is $r_{jk} = \xi_{\mathrm{F}k}/n_{\mathrm{M}}$. By convention, $r_{jk} = 0$ if a division by 0 occurs

in the formula. Each of these functions models different characteristics of how individuals enter new partnerships. A good choice of the function r_{jk} should depend on the characteristics of the population being studied, and need not be limited to those functions described here.

A.2. Derivation of pseudo-equilibrium prevalence for homogeneous SIS infection

Let π_u and π_p be the fraction of infected individuals in the unpartnered and partnered populations, respectively. Two relationships that must hold in equilibrium are derived by examining the Markov chain experienced by a 'typical' individual during a cycle of partnership formation and separation. Let $g_{\bullet,\ell}$ be the fraction of partnerships that terminate with exactly ℓ infected individuals ($\ell = 0, 1, 2$), and let $g_{i,\ell}$ be the fraction of partnerships that terminate with ℓ infected individuals, given that the partnership started with exactly i infected individuals. This notation is also included in Table 1.

First, someone infected at the end of a partnership is infected at the start of the next partnership with probability $\lambda/(\lambda + 1/d)$. For infection levels to remain stable, the following relation must therefore be true:

$$\pi_u = \frac{\lambda}{(\lambda + 1/d)}\pi_p.$$

A second relationship between π_u and π_p is established by analyzing the infection process during a partnership. To determine the $g_{\bullet,\ell}$, the values $g_{i,\ell}$ are required, as well as the probability that a partnership starts with a given number of infected partners. Assuming uniform mixing, the probability that a partnership starts with 0, 1 and 2 infected individuals is $(1-\pi_u)^2$, $2\pi_u(1-\pi_u)$, and π_u^2, respectively. The $g_{i,\ell}$ are readily determined from conditioning on the next event (infection, recovery, or separation), whose continuous rates are illustrated in Fig. 2.

The rates are readily translated into the following relationships:

$$g_{0,0} = 1, \quad g_{1,0} = \frac{1/d}{\sigma + h\beta + 1/d} + \frac{1/d}{\sigma + h\beta + 1/d}g_{2,0}, \quad g_{2,0} = \frac{2/d}{\sigma + 2/d}g_{1,0}.$$

These relationships imply that

$$g_{0,0} = 1, \quad g_{1,0} = \frac{(1/d)(\sigma + 2/d)}{(\sigma + h\phi + 1/d)(\sigma + 2/d) - 2h\beta/d},$$

$$g_{2,0} = \frac{2/d^2}{(\sigma + h\beta + 1/d)(\sigma + 2/d) - 2h\beta/d}.$$

Similarly,

$$g_{1,1} = \frac{\sigma(\sigma + 2/d)}{(\sigma + h\beta + 1/d)(\sigma + 2/d) - 2h\beta/d}, \quad g_{2,1} = \frac{2\sigma/d}{(\sigma + h\beta + 1/d)(\sigma + 2/d) - 2h\beta/d},$$

$$g_{1,2} = \frac{\sigma h\beta}{(\sigma + h\beta + 1/d)(\sigma + 2/d) - 2h\beta/d}, \quad g_{2,2} = \frac{\sigma(\sigma + h\beta + 1/d)}{(\sigma + h\beta + 1/d)(\sigma + 2/d) - 2h\beta/d}$$

By noting that

$$(\sigma + h\beta + 1/d)(\sigma + 2/d) - 2h\beta/d = \sigma^2 + \sigma h\beta + 3\sigma/d + 2/d^2,$$

it is easily verified that $g_{\mathrm{i},0} + g_{\mathrm{i},1} + g_{\mathrm{i},2} = 1$, as required. The fraction of infected individuals at the instant a partnership is broken up is then

$$\pi_{\mathrm{p}} = \frac{2g_{\bullet,2} + g_{\bullet,1}}{2},$$

where

$$g_{\bullet,2} = \sum_{i=0}^{2} p(\text{start with } i) g_{\mathrm{i},2} = \frac{2\pi_{\mathrm{u}}(1-\pi_{\mathrm{u}})\sigma h\beta + \pi_{\mathrm{u}}^2\sigma(\sigma + h\beta + 1/d)}{\sigma^2 + \sigma h\beta + 3\sigma/d + 2/d^2},$$

$$g_{\bullet,1} = \sum_{i=0}^{2} p(\text{start with } i) g_{\mathrm{i},1} = \frac{2\pi_{\mathrm{u}}(1-\pi_{\mathrm{u}})\sigma(\sigma + 2/d) + \pi_{\mathrm{u}}^2 2\sigma/d}{\sigma^2 + \sigma h\beta + 3\sigma/d + 2/d^2}.$$

Combining terms leads to the relationship in Eq. (5). Recall that $\pi_{\mathrm{u}} = \pi_p\lambda/(\lambda + 1/d)$ to obtain

$$0 = \pi_{\mathrm{u}}\left[\sigma(\sigma + 2h\beta + 2/d) - \left(\sigma^2 + \sigma h\beta + 3\sigma/d + 2/d^2\right)\left(\frac{\lambda + 1/d}{\lambda}\right)\right] - \pi_{\mathrm{u}}^2[\sigma h\beta].$$

This equation has two roots. The first is the equilibrium prevalence, $\pi_{\mathrm{u}} = 0$, and the other is the root in Eq. (6). The overall pseudo-equilibrium prevalence is then the weighted average of the prevalence in the partnered and unpartnered populations, as claimed in Eq. (7).

A.3. Derivation of basic reproduction number R_0 for simple SIS model

Let X, Y, Z be defined as in Section 2.1, and Eqs. (8)–(10). Substitute Z into the equation for Y so that

$$Z = \frac{1/d}{\sigma + 2/d}\left[\frac{\sigma}{\sigma + h\beta + 1/d}X + \frac{h\beta}{\sigma + h\beta + 1/d}[1 + Z]\right] + \frac{\sigma}{\sigma + 2/d}X,$$

$$Z = \frac{1/d}{\sigma + 2/d}\frac{h\beta}{\sigma + h\beta + 1/d} + \left[\frac{1/d}{\sigma + 2/d}\frac{\sigma}{\sigma + h\beta + 1/d} + \frac{\sigma}{\sigma + 2/d}\right]X + \frac{1/d}{\sigma + 2/d}\frac{h\beta}{\sigma + h\beta + 1/d}Z.$$

Therefore,

$$Z = \left(\frac{1/d}{\sigma + 2/d}\frac{h\beta}{\sigma + h\beta + 1/d} + \left[\frac{1/d}{\sigma + 2/d}\frac{\sigma}{\sigma + h\beta + 1/d} + \frac{\sigma}{\sigma + 2/d}\right]X\right) \Bigg/ \left(1 - \frac{1/d}{(\sigma + 2/d}\frac{h\beta}{\sigma + h\beta + 1/d)}\right)^{-1/2}.$$

Algebra indicates that

$$Z = \frac{h\beta/d + \sigma(\sigma + h\beta + 2/d)X}{(\sigma + 2/d)(\sigma + h\beta + 1/d) - h\beta/d} = \frac{h\beta/d}{(\sigma + 1/d)(\sigma + h\beta + 2/d)} + \frac{\sigma}{(\sigma + 1/d)}X.$$

Substitute Z into the equation for Y to obtain

$$Y = \frac{\sigma}{\sigma + h\beta + 1/d}[X] + \frac{h\beta}{\sigma + h\beta + 1/d}\left[1 + \frac{h\beta/d}{(\sigma + 1/d)(\sigma + h\beta + 2/d)} + \frac{\sigma}{(\sigma + 1/d)}X\right],$$

which simplifies to

$$Y = \left[\frac{\sigma}{\sigma + 1/d}\right]X + \frac{h\beta(\sigma + 2/d)}{(\sigma + 1/d)(\sigma + h\beta + 2/d)}.$$

The above expressions for Y and Z can be substituted into the equation for X to get

$$X = \frac{\lambda}{\lambda + 1/d}\left[\left[\frac{\pi_u h\beta/d}{(\sigma + 1/d)(\sigma + h\beta + 2/d)} + \frac{\pi_u \sigma}{(\sigma + 1/d)}X\right] + \left[\frac{(1 - \pi_u)\sigma}{\sigma + 1/d}X + \frac{(1 - \pi_u)h\beta(\sigma + 2/d)}{(\sigma + 1/d)(\sigma + h\beta + 2/d)}\right]\right].$$

This simplifies to the assertion of Eq. (11).

References

[1] A.L. Adams, D.C. Barth-Jones, S.E. Chick, J.S. Koopman, Simulations to evaluate HIV vaccine trial designs, Simulation 71 (4) (1998) 228.

[2] A.L. Adams, S.E. Chick, J.S. Koopman, P. Yu, GERMS: an epidemiologic simulation tool for studying geographic and social effects on infection transmission, in: P.A. Farrington, H.B. Nembhard, D. Sturrock, J. Evans (Eds.), Proceedings of the Winter Simulation Conference, Institute of Electrical and Electronics Engineers, 1999, p. 1549.

[3] M. Altmann, Susceptible-infected-removed epidemic models with dynamic partnerships, J. Math. Biol. 33 (1995) 661.

[4] M. Altmann, The deterministic limit of infectious disease models with dynamic partners, Math. Biosci. 150 (1998) 153.

[5] S.P. Blythe, C. Castillo-Chavez, J.S. Palmer, M. Cheng, Toward a unified theory of sexual mixing and pair formation, Math. Biosci. 107 (1991) 379.

[6] C. Castillo-Chavez, S. Busenberg, K. Gerow, Pair formation in structured populations, in: J.A. Goldstein, F. Kappel, W. Schappacher (Eds.), Differential equations with applications in biology, physics, and engineering, Lecture Notes In Pure And Applied Mathematics, vol. 133, Marcel Dekker, New York, 1991.

[7] O. Diekmann, K. Dietz, J.A.P. Heesterbeek, The basic reproduction ratio for sexually transmitted diseases. Theoretical considerations, Math. Biosci. 107 (1991) 325.

[8] K. Dietz, K.P. Hadeler, Biological models for sexually transmitted diseases, J. Math. Biol. 26 (1988) 1.

[9] K. Dietz, J.A.P. Heesterbeek, D.W. Tudor, The basic reproduction ratio for sexually transmitted diseases. Part 2. Effects of variable HIV infectivity, Math. Biosci. 117 (1993) 35.

[10] G.P. Garnett, R.M. Anderson, Balancing sexual partnerships in an age and activity stratified model of HIV transmission in heterosexual populations, IMA J. Math. Appl. Med. Biol. 11 (3) (1994) 161.

[11] A.C. Ghani, J. Swinton, G.P. Garnett, The role of sexual partnership networks in the epidemiology of gonorrhea, Sexually Transmitted Diseases 24 (1) (1997) 45.

[12] J.D.F. Habbema, S.J. deVlas, A.P. Plaisier, G.J. vanOortmarssen, The microsimulation approach to epidemiologic modeling of helmenthic infections with special reference so schistosomiasis, Am. J. Tropical Med. Hygiene 55 (supp. 5) (1996) 165.
[13] J.A. Jacquez, Compartmental Analysis in Biology and Medicine, 3rd Ed., BioMedware, Ann Arbor, MI, 1996.
[14] J.A. Jacquez, J.S. Koopman, C.P. Simon, L. Sattenspiel, T. Perry, Modeling and the analysis of HIV transmission: the effect of contact patterns, Math. Biosci. 92 (1988) 119.
[15] J.A. Jacquez, C.P. Simon, J.S. Koopman, Structured mixing: heterogeneous mixing by the definition of activity groups, in: C. Castillo-Chavez, (Ed.) Mathematical and Statistical Approaches to AIDS Epidemiology, Lecture Notes in Biomathematics, vol. 83, Springer, Berlin, 1989, pp. 301.
[16] J.S. Koopman, S.E. Chick, C.S. Riolo, A.L. Adams, M.L. Wilson, M.P. Becker, GERMS: A model of sexual networks and infection transmission in geographic and social space, Sexually Transmitted Diseases (2000) to appear.
[17] J.S. Koopman, J.A. Jacquez, C.P. Simon, B. Foxman, S. Pollock, D. Barth-Jones, A. Adams, G. Welch, K. Lange, The role of primary HIV infection in the spread of HIV through populations, JAIDS HR 14 (1997) 249.
[18] M. Kretzschmar, Deterministic and stochastic pair formation models for the spread of sexually transmitted diseases, J. Biol. Syst. 3 (1995) 789.
[19] M. Kretzschmar, K. Dietz, The effect of pair formation and variable infectivity on the spread of an infection without recovery, Math. Biosci. 148 (1998) 83.
[20] M. Kretzschmar, M. Morris, Measures of concurrency in networks and the spread of infectious disease, Math. Biosci. 133 (1996) 165.
[21] M. Morris, M. Kretzschmar, Concurrent partnerships and the spread of HIV, AIDS 11 (1997) 641.
[22] M. Morris, M. Kretzschmar, Concurrent partnerships and transmission dynamics in networks, Social Networks 17 (1995) 299.
[23] W.Y. Tan, X. Zhu, A Stochastic model of the HIV epidemic or heterosexual transmission involving married couples and prostitutes: I. The probabilities of HIV transmission and pair formation, Math. Comput. Modelling 24 (1996) 47.
[24] C.B.P. van der Ploeg, C. van Vliet, S.J. de Vlas, J.O. Ndinya-Achola, L. Fransen, G.J. van Oortmarssen, J.D.F. Habbema, STDSIM: a microsimulation model for decision support in STD control, Interfaces 28 (1998) 84.
[25] G. Welch, S.E. Chick, J.S. Koopman, Effect of concurrent partnerships and sex-act rate on Gonorrhea prevalence, Simulation 71 (4) (1998) 242.
[26] J.A. Yorke, H.W. Hethcote, A. Nold, Dynamics and control of the transmission of Gonorrhea, Sexually Transmitted Diseases 5 (1978) 51.
[27] S. Ross, Stochastic Processes, 2nd Ed., Wiley, New York, 1996.
[28] A.G. Frederickson, A mathematical theory of age structure in sexual populations: random mating and monogamous marriage models, Math. Biosci. 10 (1971) 117.
[29] D.D. McFarland, Comparison of alternative marriage models, in: T.N.E. Greville (Ed.), Population Dynamics, Academic Press, New York, London, 1972, p. 89.
[30] J.A. Jacquez, C.P. Simon, The stochastic SI model with recruitment and death: 1. Comparison with the closed SIS model, Math. Biosci. 117 (1993) 77.
[31] S.M. Ethier, T.G. Kurtz, Markov Processes: Characterization and Convergence, Wiley, New York, 1986.
[32] A.M. Law, W.D. Kelton, Simulation Modeling and Analysis, 2nd ed., Mc-Graw Hill, New York, 1991.