PLOS | SUBMISSION

Modeling disease spread in populations with birth, death, and concurrency

Joel C. Miller[1,2,3*] Anja C. Slim[1,4]

1 School of Mathematical Sciences, Monash University, Clayton, VIC, Australia
2 School of Biological Sciences, Monash University, Clayton, VIC, Australia
3 Institute for Disease Modeling, Bellevue, WA, USA
4 School of Earth, Atmosphere, and the Environment, Monash University, Clayton, VIC, Australia

* joel.c.miller.research@gmail.com

Abstract

The existence of sexual partnerships that overlap in time (concurrent relationships) is believed by some to be a significant contributing factor to the spread of HIV, although this is controversial. We derive an analytic model which allows us to investigate and compare disease spread in populations with and without concurrency. We can identify regions of parameter space in which its impact is negligible, and other regions in which it plays a major role. We also see that the impact of concurrency on the initial growth phase can be much larger than its impact on the equilibrium size. We see that the effect of concurrency saturates, which leads to the perhaps surprising conclusion that interventions targeting concurrency may be most effective in populations with low to moderate levels of concurrency.

Author Summary

We consider the spread of an infectious disease through a population modeled by a dynamic network with demographic turnover. We develop a stochastic model of the disease and derive governing equations that exactly predict the large population (deterministic) limit of the stochastic model. We use this to investigate the role of concurrency and find that interventions targeting concurrency may be most effective in populations with lower levels of concurrency.

Our model is not intended to be an accurate representation of any single population. Rather it is intended to give general insights for intervention design and to provide a framework which can be further specialized to particular populations.

This model is the first model to allow for analytic investigation of the impact of concurrent partnerships in a population exhibiting demographic turnover. Thus it will be useful for investigating the "concurrency hypothesis."

Introduction

The HIV epidemic has had significant impact worldwide, but it has had an especially large impact in SubSaharan Africa [1]. The reasons for this difference are many, complex, and not fully understood. One proposed factor is a larger frequency of sexual

PLOS | SUBMISSION

partnerships that overlap in time, the so-called "concurrency hypothesis" [2,3]. The concurrency hypothesis has received significant attention, but it is controversial [4–10].

The development and evaluation of potential interventions against infectious disease spread requires us to be able to predict the spread of the disease with and without the intervention in place. A crucial factor that must be considered in evaluating any potential intervention is the indirect effects of the intervention: If we directly preventing one infection from occurring, then this also prevents that individual from transmitting onwards, thus indirectly reducing exposures of other individuals. In turn their reduction in exposure reduces potential exposures of others further down the line.

Such a feedback mechanism is an inherently nonlinear effect, and so the full outcome of an intervention will depend nonlinearly on the direct effect. Doubling the intervention effort will not double the intervention impact. It may have greater or lesser impact depending on the magnitude of the indirect effect. Because of this, we often rely on mathematical models to predict the total intervention outcome. The required model complexity depends on the intervention under consideration, the biological properties of the disease, and the contact structure of the population.

Most mathematical models of disease spread ignore the fact that partnerships have duration. This is reasonable in a population with many partnerships and low transmission probabilities [11], but for a disease such as HIV, the relevant partnerships are often long-lasting and the typical number of overlapping partnerships is generally small. To be relevant for the concurrency hypothesis, we require that the mathematical model incorporate partnership duration in some form. Without it, it is impossible to consider partnerships that overlap in time. Additionally, given the time-scale of the HIV epidemic, demographic turnover plays an important role. There has not been a mathematical model which captures all of these features, and so most investigations turn to simulation, many of which are reviewed in [12].

The absence of an analytic model contributes to a large gap in our ability to understand concurrency. Analytic models make it much easier to interpret model predictions, to specifically isolate the impact of different effects, and to clarify underlying assumptions. Simulated results may be difficult to replicate by other researchers, the slow process of simulation may make parameter space impossible to investigate, and it may be difficult to identify which assumptions of the simulation are responsible for which outcomes. All of this contributes to the ongoing controversy about concurrency, and even miscommunication within the field. With an analytic model, these issues are significantly reduced.

Recent work [13–17] has shown that it is possible to make mathematical models for SI and SIR disease spread in networks with only a handful of equations. The approach has been generalized to networks with changing partnerships, but still assume a closed population. Because the HIV epidemic developed over decades, it is not appropriate to ignore individuals entering and leaving the at risk-population as they age.

In this paper we adapt the approach of [15–17] in order to accommodate "births" and "deaths" representing entry into and exit from the at-risk population. We show that the resulting equations accurately predict the outcome of simulations in large populations, and our primary focus is on using the model to investigate the role that concurrency can play in the spread of a "Susceptible–Infected" (SI) disease such as HIV.

Concurrency

A potentially important, but incompletely understood, factor in the spread of HIV and other sexually transmitted diseases is the role of "concurrent" relationships as compared to serial monogamy [2,3]. Whether its impact is significant is hotly debated [4–10].

To highlight the potential role of concurrency, we consider an individual "Alex" who has two partners "Bobbie" and "Charlie" over a period of 1 year. We demonstrate two
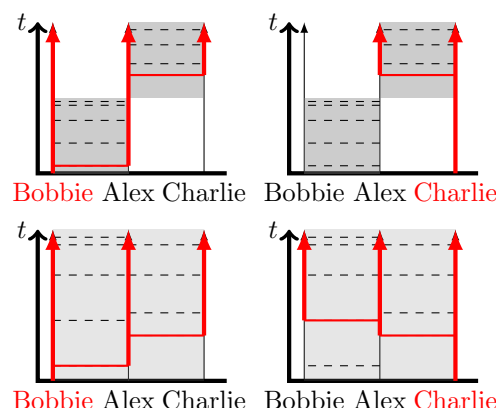
Fig 1. Sample scenarios comparing concurrent and serially monogamous relationships: Shaded regions denote the existence of a partnership. Dashed lines denote infection opportunities within the relationship that would cause infection if one individual were susceptible and the other infected. Vertical red lines denote time in which an individual is infected, and horizontal red lines denote successful transmissions. In the concurrent case, we keep exactly the same number of transmission opportunities with the gaps exactly doubled, assuming that the interaction rate within each partnership is half that of the serial case. Concurrency can speed up onward transmission and provide additional transmission routes.

potential partnership arrangements and explore their impact on disease transmission in Fig. 1.

In the first scenario, the partnership with Bobbie lasts for six months and is replaced by a six month partnership with Charlie (the serial case), while in the second scenario both partnerships overlap for the entire period (the concurrent case). In the concurrent case, the intervals between transmitting events are doubled.

Let us focus first on Alex's risk of infection if Bobbie is initially infected. In both cases, Alex has the same probability of becoming infected, though infection would occur sooner in the serial case. In contrast, if Charlie is initially infected, Alex would become infected sooner in the concurrent case (but again with the same overall probability). So the two models appear to give similar outcomes from the point of view of the cumulative risk of infection to Alex.

For onward transmission, however, we see the potential for large differences. If Bobbie is initially infected, transmission from Bobbie to Alex to Charlie will tend to happen faster in the concurrent case because there is no built-in delay while waiting for the partners to change.[1] In turn this allows Charlie to begin transmitting to other partners earlier. In addition, if Charlie is initially infected, transmission from Charlie to Alex to Bobbie is possible only in the concurrent case.

It is important to recognize that the "concurrency hypothesis" does not suggest that there is a difference in the total risk of infection of an individual who has multiple concurrent partners compared to one who has the same number of partners sequentially. Instead, it increases the risk to the partners.

As a general rule we anticipate that concurrency will increase the spread of disease through two mechanisms:

1. by allowing the disease to trace transmission routes faster, and

---

[1] However, the probability of a Bobbie to Alex to Charlie transmission chain is slightly reduced because some of the interactions between Alex and Charlie will have already happened by the time Alex is infected.

PLOS | SUBMISSION
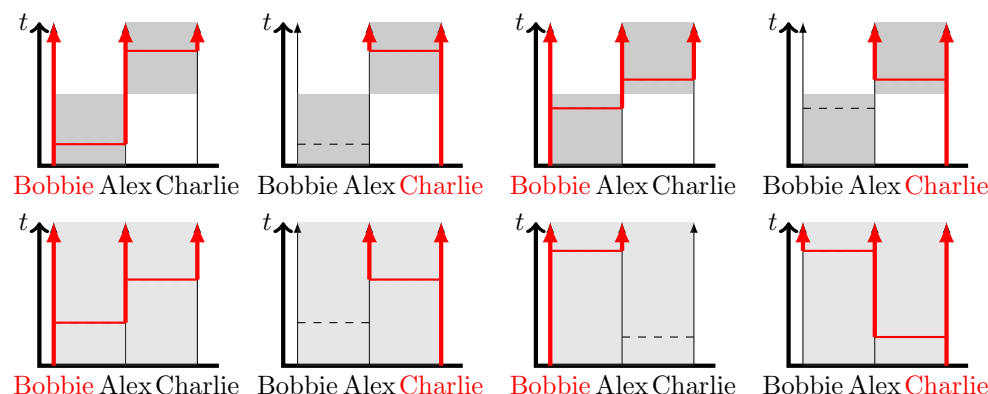
Fig 2. Scenarios for low transmission rates: In the low transmission rate limit where we can assume at most one transmission happens per partnership, the additional transmission paths and faster transmissions from concurrent partnerships do not play a significant role. There are two cases of interest: The first two columns show the case in which the Bobbie-Alex partnership transmits relatively earlier in the partnership than the Alex-Charlie partnership. The last two columns show the Alex–Charlie partnership transmitting faster.

2. by providing additional transmission routes. 84

In the low transmission rate limit, for which the possibility of transmitting multiple 85
times is negligible, we can assume that each partnership either transmits once or not at 86
all. In this case, Fig. 2 suggests that the difference between concurrent and serial 87
relationships is likely to be insignificant. 88

Conceptually we can explain this by imagining that we only observe transmission 89
events. Anything that influences the disease spread must be detectable by our 90
observations. Conversely, if we cannot detect something by observing disease 91
transmission events, then it cannot matter to the disease spread. In particular, if no 92
partnerships transmit more than once, then we cannot tell whether an individual who is 93
transmitting is doing so to one of several randomly selected current partners or if it is 94
transmitting to its current partner who is randomly selected from the potential partners. 95

Materials and Methods 96

Our goal is twofold: 97

1. To demonstrate an analytic model for disease spread in a dynamic population with 98
   concurrency and show that it accurately predicts simulations in large populations. 99

2. To use the model developed to provide guidance about the role of concurrent 100
   relationships in simplified scenarios. 101

In this section we introduce our stochastic population and disease model, state the 102
governing equations for the large population limit (the equations themselves are derived 103
in the Supporting Information), and demonstrate that they accurately reproduce 104
simulated epidemics in large populations. 105

Part of the reason for the debate about concurrency's effect can be traced to the 106
large parameter space and the amount of time it takes to simulate for different 107
parameters. Simulations have studied only a relatively small portion of parameter space, 108
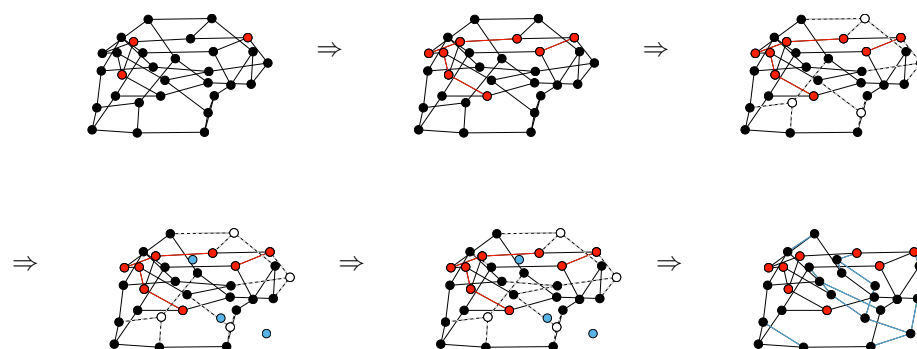
Fig 3. Sequence of events in each time step. We begin with a network with some infected individuals (red). Then infected individuals transmit to some partners (red edges). Then some individuals leave the population (white). Other individuals are born (blue). Then some edges break (dashed). Finally edges are added so that the new individuals, the individuals whose partners left, and the individuals whose edges broke all return to their target number of partners. The sequence then repeats.

and so some of the disagreement revolves around intuition about how parameter changes would affect outcomes rather than quantitative statements. Our analytic model will allow for a much more efficient exploration of parameter space.

Population/Disease Assumptions

We will use a discrete-time model. Each time step is broken down into sub-intervals in which the actions occur in a specific order, shown in Fig. 3. At the start of a time step, each partnership connecting a susceptible and infected individual transmits with probability τ. Then some individuals leave the population, each independently with probability μ. Next μN new individuals arrive where N is the typical population size. Then each remaining partnership ends independently with probability η. Finally new partnerships are formed until all individuals reach their target number of partners. These assumptions are similar to those of [18, 19].

We will present an analytic model that captures the deterministic limit of these assumptions. We have chosen a discrete-time model because the complementary simulations are easier in a discrete-time framework. In a large enough simulated population with a large enough time step, many partnerships end and many new individual enter at each time step. Thus an individual who needs a new partner has many choices. In a continuous-time model or with a small time step (relative to population size), it would be more difficult for an individual who has recently ended a partnership to find a new partner.

For simulations, we choose the time step of the discrete-time framework to balance competing interests. We want a small time step so that μ, η, and τ are small (at leading order they are proportional to the time step). When they are not small, the somewhat arbitrary order of events we have assumed will begin to impact outcomes. However, too small of a time step will mean that few partnerships end in a time step. This makes it difficult for nodes to immediately find new partners. As the population size increases this becomes less of a problem so smaller time-steps become feasible. However, the computational effort becomes greater. This is not an issue in the analytic model, and so a continuous-time model is not difficult to consider. The predictions from the

continuous-time and discrete-time models will be the same for small parameters. 138

Governing equations 139

We derive the governing equations in the Supporting Information. These equations are 140
based on the "Edge-based Compartmental Modeling" approach of [13, 15, 16]. They are 141
low-dimensional, but are significantly more involved than even the dynamic network 142
models presented in [16]. This is because the age of an individual gives us some 143
information about that individual's status. In addition the age of a partnership gives 144
information about the age of the partner. Thus in calculating the risk an individual has 145
from its partners, we need to account for the probability the partner has a given status, 146
which depends on the age of the partner, which in turn depends on the age of the edge, 147
which itself is dependent on the age of the individual. To sort out the dependencies, age 148
of individual and age of partnership are needed as independent variables. 149

Table 1 defines the variables and parameters of the model. 150

151

The model parameters are the death probability μ, the transmission probability τ, the partnership change probability η, and the initial fraction infected ρ. We seek the susceptible and infected fractions of the population S and I. The governing equations are

$$S(t) = \mu \sum_{a_u=0}^{\infty} (1-\mu)^{a_u} s(t, a_u)$$

$$s(t, a_u) = \begin{cases} \psi(\Theta(t, a_u)) & a_u < t \\ (1-\rho)\psi(\Theta(t, t)) & a_u \geq t \end{cases}$$

$$I(t) = 1 - S(t)$$

$$\Theta(t, 0) = 1$$

$$\Theta(0, a_u) = 1$$

$$\Theta(t, a_u) = \Theta(t-1, a_u-1) - \tau\Phi_I(t-1, a_u-1) \qquad t, a_u \geq 1$$

$$\Phi_I(t, a_u) = \Theta(t, a_u) - \Phi_S(t, a_u)$$

$$\Phi_S(t, a_u) = (1-p_b)^{a_u}\phi_S(t, a_u, a_u) + p_b \sum_{a_e=0}^{a_u-1} (1-p_b)^{a_e}\phi_S(t, a_u, a_e)$$

$$\phi_S(t, a_u, a_e) = \Theta(t - a_e, a_u - a_e)\chi(t, a_e)$$

$$\chi(t, a_e) = \begin{cases} (1-\rho)\frac{\psi'(\Theta(t,t))}{\langle K \rangle} & a_e \geq t \\[3mm] (1-P_e)\frac{\psi'(\Theta(t,a_e))}{\langle K \rangle} & \\ \quad + P_e\mu \sum_{a_v=a_e+1}^{t-1} (1-\mu)^{a_v-a_e-1}\Theta(t-a_e, a_v-a_e)\frac{\psi'(\Theta(t,a_v))}{\langle K \rangle} & a_e < t \\ \quad + P_e(1-\rho)\Theta(t-a_e, t-a_e)\frac{\psi'(\Theta(t,t))}{\langle K \rangle}(1-\mu)^{t-a_e-1} & \end{cases}$$

where $p_b = 1 - (1-\eta)(1-\mu)$ is the probability an individual's partnership ends in a 152
time step either because of partnership change or the partner leaving the population, 153
and $\psi(x) = \sum_k P(k)x^k$ is the probability generating function of the degree distribution. 154
We derive this model and a continuous-time differential equations version in the 155
Supporting Information. 156

Variable	Description
t	Time
u	The test individual
a_u	The age of test individual u, measured so that $a_u = 0$ in the first time step after u is born.
a_e	Age of an edge of interest, measured so that $a_e = 0$ in the first time step after the edge forms.
$S(t)$	The proportion of the population that is susceptible at time t, equivalently the probability a randomly selected individual is susceptible at time t, or equivalently the probability a test individual is susceptible at time t.
$I(t)$	$1 - S(t)$: The proportion infected at time t.
$s(t, a_u)$	The probability a test individual of age a_u is susceptible at time t.
$\Theta(t, a_u)$	The probability a stub belonging to u has not transmitted infection to it from a partner by the start of time step t.
$\Phi_S(t, a_u)$	As for Θ, but with the additional requirement that at the start of time step t the partner is susceptible.
$\Phi_I(t, a_u)$	As for Θ (no partner has transmitted to u through the stub), but with the additional requirement that at the start of time step t the partner is infected.
$\phi_S(t, a_u, a_e)$	The probability that a stub belonging to an age a_u individual has not transmitted infection to it by time t, is connected to a susceptible neighbor, and the current partnership has age a_e.
$\chi(t, a_e)$	The probability that an age a_e edge of a test individual connects to a susceptible individual.
$\psi(x)$	$\sum_{k=0}^{\infty} P(k) x^k$: The probability generating function of the degree distribution.
μ	The probability a random individual will leave the population in a given time step.
τ	The transmission probability in a time step.
η	The probability a partnership will end in a time step.
p_b	$1 - (1 - \mu)(1 - \eta)$: The probability that a test individual's partnership will end (break) either because the partnership ends naturally η, or the partner leaves the population μ. It does not include the possibility of the test individual leaving.
P_e	$\frac{(1-\mu)(\eta+\mu-\eta\mu)}{(1-\mu)(\eta+\mu-\eta\mu)+\mu}$: The probability that a newly formed partnership will be with a previously existing individual.
ρ	The proportion of the population randomly infected at $t = 0$.
b	The number of individuals entering the population each time step (assumed to be constant).
N	b/μ: The average population size.

Table 1. The variables for our simulations and equations.
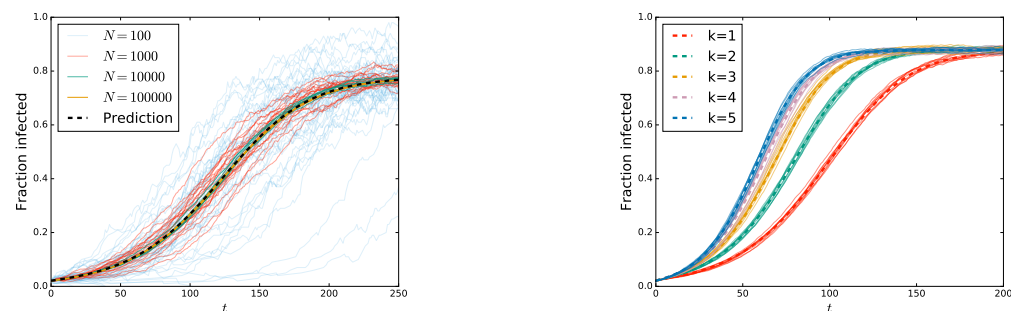
Fig 4. Comparison of dynamics from analytic prediction and stochastic simulation: (Left) We consider stochastic epidemic simulation in populations of average size $N = 10^2$, 10^3, 10^4 and 10^5 with $k = 3$, $\eta = 0.2/3$, $\tau = 0.05/3$, $\rho = 0.02$, and $\mu = 0.01$. As N increases, the simulations converge to the prediction. (Right) For $\eta_1 = 0.1$, $\tau_1 = 0.1$, $\rho = 0.02$, $\mu = 0.01$ and average population size $N = 10000$ we compare predictions and simulations for different values of k, using $\tau = \tau_1/k$ and $\eta = \eta_1/k$. We find excellent agreement.

Comparison with simulation

To demonstrate the accuracy of the equations and provide some insight into model predictions, we compare predicted dynamics (solutions of the system of equations) with stochastically simulated epidemics under different conditions in Figs. 4 and 5.

We start Fig. 4 with a plot showing that as the population size increases, the simulations converge to the predicted dynamics for one set of parameters.

The second plot of Fig. 4 considers populations in which all individuals have k partners for $k = 1, 2, 3, 4$, and 5 with $\tau = \tau_1/k$ and $\eta = \eta_1/k$ for various values of τ_1 and η_1. Simulations and predictions are a good match for different values of k. Interestingly we see that for different values of k the equilibrium level does not vary much, but the early growth rate does.

None of the populations in Fig. 4 have heterogeneous degree. Figure 5 looks at disease spread in populations with heterogeneous degrees, again showing excellent agreement between stochastic simulation and predictions.

From the first plot in Fig. 5 we infer that for a given population, the initial proportion infected does not influence the final state. Interestingly, we see that it is possible for the disease to initially grow and then decay for some region of ρ. This is because in this region the number of infected high degree nodes increases quickly from the initial condition while the number of infected low degree nodes decreases slowly. At long time, reach a steady state, with the decrease of low degree infections outweighing the increase in high degree infections.

From the second plot of Fig. 5, we infer that increasing the population turnover rate decreases the proportion infected. This is not particularly surprising as it implies that infected individuals leave the population sooner, having had less opportunity to cause further infections.

Our main conclusion from Figs. 4 and 7 is that the equations accurately predict the dynamics of simulations regardless of the parameters used. The numerical solution of these equations is much faster than simulation.

Results & Discussion

The previous section showed that the analytic equations provide accurate predictions of the disease dynamics in the large-population limit. We now specifically explore the
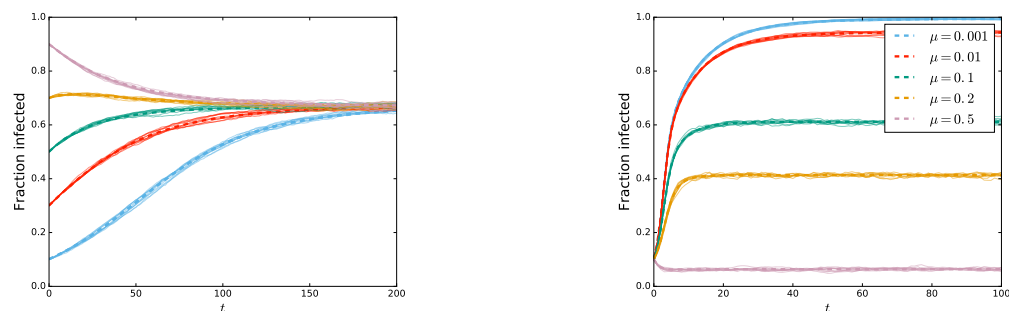
Fig 5. Comparison of dynamics from analytic prediction and stochastic simulation in populations with heterogeneous degree: (Left) Disease spread in populations with average size $N = 10000$ and degree probabilities $P(2) = P(7) = 1/2$. The parameters are $\tau = 0.01$, $\eta = 0.005$, and $\mu = 0.01$. The initial fraction infected varies. (Right) Disease spread in populations with average size $N = 10000$ and degree probabilities $P(1) = 1/2$, $P(10) = 1/3$, and $P(20) = 1/6$. The parameter μ varies between populations. The remaining parameters are $\eta = 0.05$ and $\tau = 0.1$. In both plots the dashed curves are predictions and thin solid curves are stochastic simulations.

predictions the model makes if we change the amount of concurrency in the population. The second plot of figure 4 hints at our results. We generally see relatively little impact of concurrency on the proportion infected at equilibrium, while we do see an impact on the early growth.

Although the equations allow different individuals to have a different number of partners, for the remainder of this paper we focus on populations in which all individuals have the same number of partners. In particular, this eliminates the need to consider how per-partnership transmission rates may depend on the individual's number of partners. This allows us to separate the effect of concurrency from the effect of some individuals having more frequent sexual activity than others.

So for our purposes, $\psi(x) = x^k$, $\langle K \rangle = k$ for some fixed value k. As a base-case, we consider serial-monogamy where each individual has a single partner ($k = 1$). Transmission occurs in a time-step with probability τ_1 and partnerships end with probability η_1. Individuals leave the population with probability μ. We compare this with homogeneous populations having concurrency.

We assume that each population is arranged such that the number of partners an individual has over a long period of time is the same. This implies that $\eta = \eta_1/k$ so that the k partnerships each lasts k times as long. Similarly we assume that the expected number of transmissions an infected individual would cause in a time step is the same. This requires $\tau = \tau_1/k$.

Because of how we have set up our populations, we can rule out many causes for the differences we see. We know that the effects we observe are not explained by within-population heterogeneity in degree, within-population heterogeneity in sexual activity rates, between-population differences in typical life-time number of partners, or between-population differences in the number of transmissions an individual causes per time step. All of these effects have been removed. The only remaining difference between the populations is the number of concurrent partnerships.

Impact on equilibrium size

We see in Fig. 6 that the impact of concurrency on the equilibrium epidemic size can be significant, but that the effect of increasing k saturates quickly. Little changes once $k > 2$. This suggests that reducing concurrency may play an important role in reducing
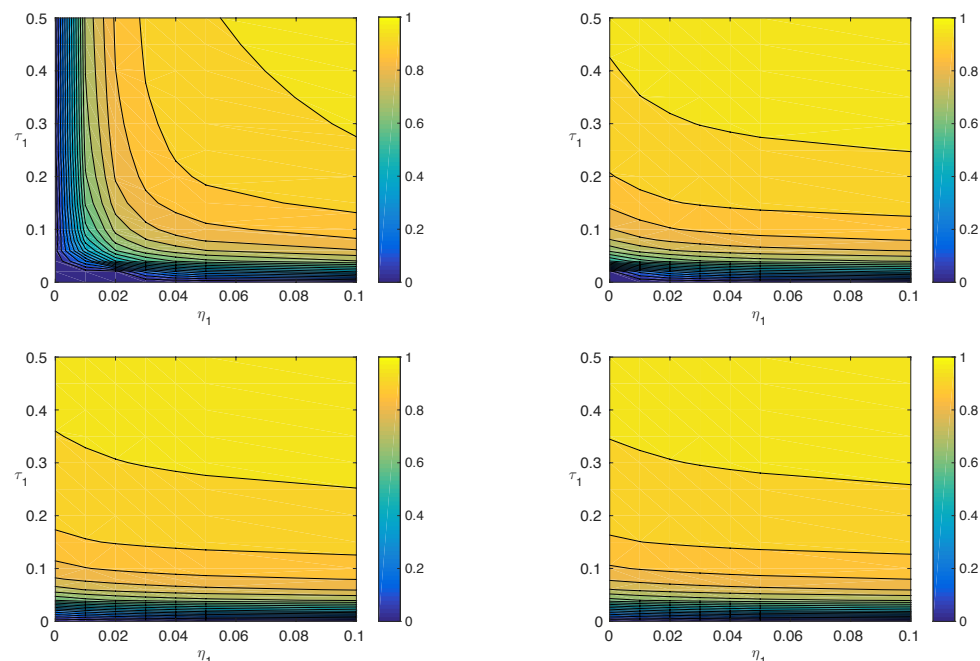
Fig 6. Comparison of equilibrium sizes for different values of k: The equilibrium fraction infected for $\mu = 0.01$ and different η_1 and τ_1. We use the same axes η_1 and τ_1, taking $\eta = \eta_1/k$ and $\tau = \tau_1/k$. As k increases, the figures quickly converge. The effect of concurrency on the equilibrium size quickly saturates.

the epidemic size, but that if there is significant clustering, a small or moderate change is unlikely to have a significant impact.

From the perspective of intervention design, this suggests that interventions which aim to reduce concurrency may be most effective in populations with lower rates of concurrency already. We can think of this analagously with interventions targetting $\mathcal{R}_0$ in an SIR disease. Reducing $\mathcal{R}_0$ from a large number to a moderate number has little effect on the predicted final size, while a reduction from a moderate number to a small number (even if still above 1) would have a much larger impact.

We note that close to the epidemic threshold the impact of any parameter change (including amount of concurrency present) can be significant.

Impact on early growth

Although Fig. 6 appears to suggest concurrency is unimportant, this is only for the equilibrium size. We saw in Fig. 4 that concurrency can play a role when we consider the early growth rate of the disease even if the equilibrium size is unaffected. That is, concurrency facilitates early spread of a disease. We demonstrate this more generally in Fig. 7. Although the impact of increasing k will eventually saturate, it takes much longer before saturating for early growth than for equilibrium size.

This makes sense: if a disease invades a population with serial monogamy, then no matter how infectious it is, it must wait until a partnership ends before it can spread further. However, if there is concurrency, this restriction vanishes, and the spread is not restricted by the partnership time scale. If there are only a few partners, then transmissions lead to localized depletion of susceptibles around an infected individual, stifling further transmissions. As the number of partners increases with a corresponding decrease in transmission rate, the number of transmission opportunities remains the
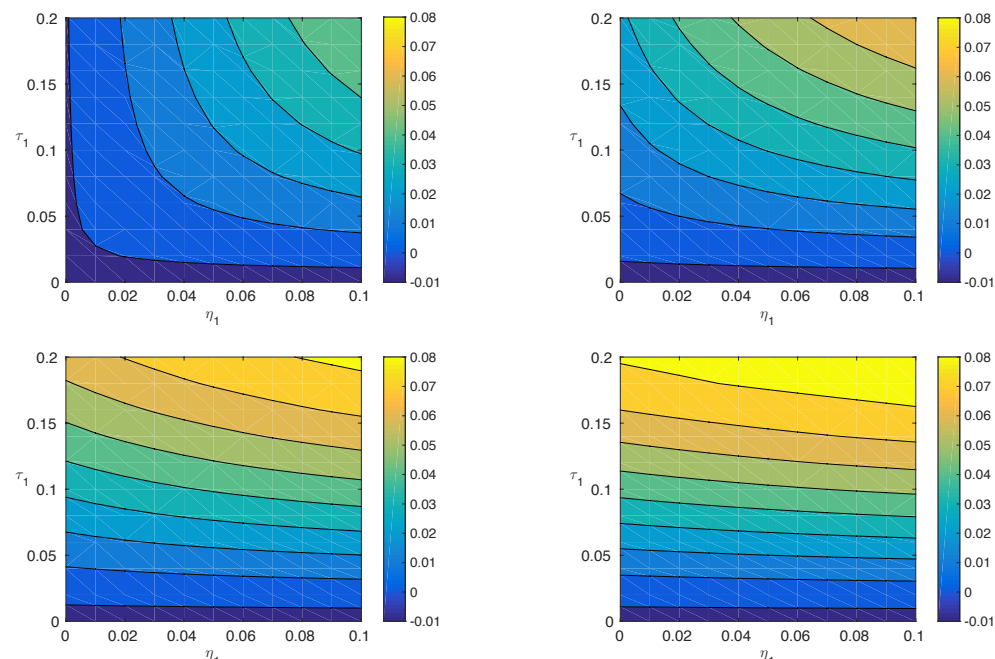
Fig 7. Comparison of early growth with different levels of concurrency:
Contour plots of $\Delta I / I$ once the dynamics have entered the early exponential phase.
Although eventually the change saturates with larger k, it does not saturate nearly as
quickly as does the equilibrium size.

same but the transmissions go to more individuals.

Caveats

There are a number of caveats of our study that must be highlighted to avoid
overinterpreting these results. To simply the model we present, we have neglected many
effects. Many of these can be easily incorporated into more sophisticated versions of the
model.

- **Acute Phase:** Before mounting an immune response, an individual's viral load is
  several orders of magnitude larger than after the immune response develops.
  During this early phase infectiousness is dramatically increased [20, 21]. If the
  individual has multiple partnerships, then many more infections can happen in
  this phase than would be seen if the individual were only in contact with its
  infector.

- **Heterogeneous degree:** Some people have many more partnerships than
  others [22]. They generally become infected sooner, and in turn transmit to more
  individuals. Even if many individuals do not engage in concurrent relationships, if
  there are some who have many partners, the effects may still be present. This
  provides an opportunity to reduce disease transmission through an intervention
  that encourages those without concurrent relationships to select partners without
  concurrent relationships.

- **Temporal behavior changes:** If the disease dynamics are driven by individuals
  having periodic high-risk episodes between long-term relationships, then the

assumptions of this model are invalid. To correct this, the model must be adapted   264
to allow for periods of high risk behavior, for example when a partnership ends.   265

- **Age structure:** If there is age-structure in the contact patterns, different effects   266
  may be seen. For example, we might think of the younger cohort as a population   267
  which has not yet been invaded by infection. In this case, the results about disease   268
  invading may be more relevant than the results about population size. Here   269
  reducing concurrency could be expected to play an important role in slowing the   270
  invasion of this younger cohort.   271

- **Coital dilution:** We have assumed that the transmission rate scales such that   272
  individuals have effectively the same number of sexual acts regardless of their   273
  number of partners. This allows us to isolate the effect of concurrency from the   274
  effect of frequency of sexual acts. However, if concurrent relationships are   275
  associated with more frequent sexual acts, then the conclusions we reach here may   276
  not be valid. To correct for this, we need to appropriately weight the transmission   277
  rates based on the number of concurrent relationships each partner has.   278

Conclusion   279

We have derived an analytic model which accurately reproduces simulated SI epidemics   280
in a population with concurrent relationships and demographic turnover. We use this   281
model to isolate the role of concurrency in the spread of a disease such as HIV. In   282
isolating concurrency, we consciously choose to neglect a number of other important   283
effects.   284

Although the model is highly simplistic, it gives insight into the role of concurrency.   285
We see first that the impact of concurrency on the equilibrium size of SI epidemics   286
appears to saturate quite quickly. Consequently we might expect that interventions   287
targeting concurrency will have little impact unless they come close to eliminating   288
concurrent relationships.   289

However, we see a much larger role for concurrency in determining the early growth   290
rate. As concurrency increases, the early growth is increased, and the impact of   291
concurrency saturates much slower than for the epidemic size. So for slowing the initial   292
invasion of an epidemic, reducing concurrency is likely to have an important impact.   293

Our observations that the impact of concurrency saturates suggests care may be   294
needed before diverting resources from some intervention to an intervention aimed at   295
reducing concurrency. A more realistic model is needed, and some consideration will   296
need to be made for whether the population is already in equilibrium or whether the   297
epidemic is still growing.   298

The model presented here is intended as a framework for developing more detailed   299
models. Our goal in introducing this model has been to provide this framework and   300
clearly demonstrate that it is possible to use analytic models to explore disease spread   301
in populations with concurrent relationships with demographic turnover. The   302
predictions our model has provided are true for the simplistic assumptions made. More   303
careful models will be needed to identify conditions under which interventions targeting   304
concurrency will be effective. These models will need to incorporate additional effects   305
such as the acute phase of infection and more realistic information about degree   306
distributions and correlations.   307

PLOS | SUBMISSION

Supporting Information   308

Our primary goal in the Supporting Information is to derive the governing equations.   309
For simplicity, we have ignored many effects. Although in the main text we assume all   310
individuals have the same number of concurrent partners, in our derivation here we   311
allow different individuals within the same population to have a different number of   312
concurrent partners as long as all partnerships have the same transmission probability   313
and typical duration. These assumptions could be modified, and a number of other   314
complexities added to the model we develop here, but we do not attempt this now.   315

Stochastic population and disease model   316

We now describe the stochastic rules we assume govern the population and disease   317
dynamics. We use a discrete-time model. We begin with the population dynamics in the   318
absence of disease. At each time step, $N\mu$ individuals enter the population, and each   319
individual has probability μ to independently leave the population. This leads to an   320
equilibrium population size of N, but with variation around this value.   321

Each individual u has a constant number of partners k_u which is assigned   322
independently to u when u enters the population. $P(k)$ gives the probability that   323
$k_u = k$. We think of u as having k_u "stubs" (also called "binding sites" by [18,19]). The   324
stubs pair with stubs from other individuals to form partnerships. When a partnership   325
ends the two newly freed stubs join with other free stubs to form new partnerships. We   326
assume that individuals immediately replace their partners so that at the start of each   327
time step all individuals have a full set of partners.[2]   328

We are interested in the epidemic timescale, which is longer than the individual's   329
active period. So we must include "birth" and "death" or equivalently immigration and   330
emigration.   331

We begin with a fraction ρ of the population randomly infected. In each time step,   332
multiple events can happen. Since the order of events can matter (a partnership cannot   333
transmit after it ends), we provide a consistent order, shown in figure 3. First, infected   334
individuals transmit to their susceptible partners independently with probability τ.   335
Second, individuals may "die" (or leave the population) independently with probability   336
μ. Third, μN new individuals are added to the population (so the average number   337
present is N) and assigned stubs. Fourth, each remaining partnership breaks with   338
probability η. Finally, the unpaired stubs form new partnerships, subject to the   339
constraint that old partnerships are not reformed and individuals do not join to   340
themselves. In simulations, these constraints are occasionally not satisfied, in which case   341
the corresponding individuals wait a time step before attempting to reform edges. In a   342
large population, the impact of this failure is negligible, and for our analytic equations   343
below, we can assume that they are satisfied.   344

Equation Derivation   345

Preliminaries   346

It will be useful to define the function   347

$$\psi(x) = \sum_k P(k)x^k$$

to be the probability generating function of the degree distribution. It has some   348
important properties: $\psi(1) = \sum P(k)1^k = 1$, $\psi'(1) = \sum kP(k)1^{k-1} = \langle K \rangle$ where $\langle X \rangle$ 349
denotes the mean of the random variable x.   350

---

[2]In [16] there is discussion of how to include more complicated partnership dynamics.

We now provide a derivation of the deterministic equations governing the large population limit of our model. The derivation is based on [16]. We review the concept of a "test individual" (effectively equivalent to the *cavity state* of [23]). We start with the assumption that the population-scale dynamics are deterministic in the large population limit. A direct consequence of this assumption is the observation that the probability a randomly selected individual has a given status equals the proportion of the population with that status. 351 352 353 354 355 356 357

Calculating the probability a random individual has a given status turns out to be simpler than calculating the number of individuals in each state. An important observation is that the probability a single randomly chosen individual u has a given status is not affected if we prevent it from infecting any other individuals[3] If we prevent u from transmitting to its neighbors, then the status of its neighbors become independent of one another, but this does not alter the probability that u is susceptible. 358 359 360 361 362 363

We define a *test individual* to be an individual u chosen uniformly at random from the population and prevented from transmitting infection. We have the following sequence of questions which have identical answers if the dynamics are deterministic: Given the initial conditions, 364 365 366 367

1. What fraction of individuals are susceptible or infected at time t? 368

2. What is the probability a random individual is susceptible or infected at time t? 369

3. What is the probability a randomly chosen test individual is susceptible or infected at time t? 370 371

In our derivation we assume that a fraction ρ is randomly infected at $t = 0$. As long as $\rho N \gg 1$, our equations are appropriate. 372 373

To begin our calculations we start with $\Theta(t, a_u)$, the probability that a stub belonging to an age a_u test individual u has never been involved in a transmission to u. Once we know that, then the probability a test individual of age a_u and k_u partners is susceptible at time t is $\Theta(t, a_u)^{k_u}$. Averaging this over the entire population of age a_u individuals the probability an age a_u individual is susceptible is 374 375 376 377 378

$$s(t, a_u) = \begin{cases} \psi(\Theta(t, a_u)) & a_u < t \\ (1 - \rho)\psi(\Theta(t, a_u)) & a_u \geq t \end{cases}$$

where we recall $\sum_k P(k)x^k = \psi(x)$, and the $1 - \rho$ factor in the second term accounts for the fact that the individual would be infected at $t = 0$ with probability ρ. 379 380

The probability that a random individual present at time t has age a_u is $\mu(1 - \mu)^{a_u}$. The fraction susceptible is thus 381 382

$$S(t) = \mu \sum_{a_u=0}^{\infty} (1 - \mu)^{a_u} s(t, a_u)$$

The probability of being infected is 383

$$I(t) = 1 - S(t)$$

The focus of our calculations is on determining $\Theta(t, a_u)$. As a boundary condition we have 384 385

$$\Theta(t, 0) = 1$$

---

[3]Although it is not necessary here, it may be helpful to recognize that that the assumption the stochastic process exhibits deterministic population-scale dynamics means that a change of out come for a vanishingly small fraction of events does not alter the population-scale dynamics.

stating that when an individual is first introduced, it has not yet received any infection. Similarly we have the initial condition

$$\Theta(0, a_u) = 1$$

as well, stating that prior to the disease introduction, no transmissions have occurred. The change in Θ in a time step is $-\tau\Phi_I$ where Φ_I is the probability that the stub has not previously brought infection to u and connects to an infected partner at the start of the time step. So we have

$$\Theta(t, a_u) = \Theta(t-1, a_u-1) - \tau\Phi_I(t-1, a_u-1)$$

However to do this calculation we require $\Phi_I(t, a_u)$ which is still unknown. We can shift our unknown from Φ_I to Φ_S (the probability the stub has not transmitted to u and currently connects to a susceptible partner) by using

$$\Phi_I = \Theta - \Phi_S \,.$$

As in calculating S, to calculate Φ_S, we turn it into a sum. The probability that a partnership created when u joined still exists is $(1-p_b)^{a_u}$. The probability that a partnership has some smaller age a_e is $p_b(1-p_b)^{a_e}$. So

$$\Phi_S = (1-p_b)^{a_u} \phi_S(t, a_u, a_u) + p_b \sum_{a_e=0}^{a_u-1} (1-p_b)^{a_e} \phi_S(t, a_u, a_e)$$

where $\phi_S(t, a_u, a_e)$ is the probability that a stub belonging to an age a_u individual that is part of an age a_e edge has not transmitted by time t and p_b is the probability that a stub is freed to find a new partnership (either by death of the partner, or termination of the partnership). The one term outside the sum represents the fact that when the individual first enters the population the stub definitely forms a partnership.

We now find $\phi_s(t, a_u, a_e)$. If the edge formed when u was born ($a_e = a_u$) then this is simply the probability the partner v is susceptible given that v has an age a_u partnership with u, which we denote $\chi(t, a_u)$. However, if the edge formed after u was born ($a_e < a_u$), then $\phi_S(t, a_u, a_e)$ is the probability $\Theta(t - a_e, a_u - a_e)$ that the stub was not responsible for transmitting infection to individual u prior to the current partnership forming times $\chi(t, a_e)$. As $\Theta(t - a_u, 0) = 1$ these coincide when $a_u = a_e$, so we can write

$$\phi_S(t, a_u, a_e) = \Theta(t - a_e, a_u - a_e)\chi(t, a_e)$$

We now find $\chi(t, a_e)$ similarly to $s(t, a_u)$. It is

$$\chi(t, a_e) = \sum_{A_v=a_e}^{\infty} P(a_v = A_v | a_e) P(v \text{ susceptible} | a_v = A_v) \,.$$

If $a_e \geq t$, then we know that v was born either when the disease was introduced or earlier. Thus no previous partnership could have transmitted to v. if we assume $a_v = A_v \geq a_e$, then the probability v is susceptible is the probability that it escaped infection when the disease was introduced $1 - \rho$ times the probability that it has not been infected by any other partners. Because of how v is selected (it is u's partner), v is likely to have a higher degree than a randomly selected individual. The probability v has degree $k_v = k$ is $kP(k)/\langle K \rangle$. So the probability v is susceptible given A_v is $(1-\rho)\sum_k [kP(k)/\langle K \rangle]\Theta(t, A_v)^{k-1} = (1-\rho)\psi'(\Theta(t, A_v))/\langle K \rangle$. Thus for $a_e \geq t$ we have

$$\chi(t, a_e) = \sum_{A_v=a_e}^{\infty} P(a_v = A_v | a_e) P(v \text{ susceptible} | a_v = A_v)$$

$$= \sum_{a_v=a_e}^{\infty} \mu(1-\mu)^{a_v-a_e}(1-\rho)\frac{\psi'(\Theta(t, a_v))}{\langle K \rangle} \qquad a_e \geq t$$

For $a_e < t$ there are three important cases to consider based on whether the partner was born before or at the same time that the partnership was formed and whether the partner was born before or after the disease was introduced.

- If the edge formed when v was born then $A_v = a_e$ for which
  $P(a_v = a_e|a_e) = (1 - P_e)$ and
  $P(v \text{ susceptible}|a_v = a_e) = \sum_k [kP(k)/\langle K \rangle]\Theta(t, a_e)^{k-1} = \psi'(\Theta(t, a_e))/\langle K \rangle$,
  which measures the probability that another partner of v has not transmitted to v.

- If v was born before the edge formed but after the disease was introduced then
  $a_e < A_v < t$ and $P(a_v = A_v|a_e) = P_e\mu(1 - \mu)^{A_v - a_e - 1}$. Although u has not transmitted to v, it is possible that a previous partner of v that was eventually replaced by u did. Thus the probability v is susceptible is
  $\Theta(t - a_e, A_v - a_e)\psi'(\Theta(t, A_v))$.

- If v was born before the disease was introduced, then $A_v \geq t$. We again have
  $P(a_v = A_v|a_e) = P_e\mu(1 - \mu)^{A_v - a_e - 1}$, but there is an extra factor of $1 - \rho$ in the probability v is susceptible.
  $P(v \text{ susceptible}|A_v) = (1 - \rho)\Theta(t - a_e, A_v - a_e)\psi'(\Theta(t, A_v))$.

So for $a_e < t$ we have

$$
\begin{aligned}
\chi(t, a_e) &= \sum_{A_v = a_e}^{\infty} P(a_v = A_v|a_e)P(v \text{ susceptible}|a_v) \\
&= P(a_v = a_e|a_e)P(v \text{ susceptible}|a_v = a_e) \\
&\quad + \sum_{A_v = a_e + 1}^{t-1} P(a_v = A_v|a_e)P(v \text{ susceptible}|a_v = A_v) \\
&\quad + \sum_{A_v = t}^{\infty} P(a_v = A_v|a_e)P(v \text{ susceptible}|a_v = A_v) \\
&= (1 - P_e)\frac{\psi'(\Theta(t, a_e))}{\langle K \rangle} \\
&\quad + P_e\mu \sum_{A_v = a_e + 1}^{t-1} (1 - \mu)^{A_v - a_e - 1}\Theta(t - a_e, A_v - a_e)\frac{\psi'(\Theta(t, A_v))}{\langle K \rangle} \\
&\quad + P_e\mu(1 - \rho) \sum_{A_v = t}^{\infty} (1 - \mu)^{A_v - a_e - 1}\Theta(t - a_e, A_v - a_e)\frac{\psi'(\Theta(t, A_v))}{\langle K \rangle}
\end{aligned}
$$

Simplification for $a_u > t$

We claim that the value of $\Theta(t, a_u)$ is the same for all $a_u \geq t$. This follows from the fact that at $t = 0$ all the values are 1. By inspecting the equations for the evolution of Θ, we see that if we assume $\Theta(t, a_u)$ is the same for all $a_u \geq t$, then the change in Θ is also the same. Thus we can assume $\Theta(t, a_u) = \Theta(t, t)$ if $a_u > t$. This argument would break down if partnership formation were affected by age differences.

Among the resulting simplifications is the observation that for $a_e \geq t$, the expression for $\chi(t, a_e)$ simplifies to $(1 - \rho)\psi'(\Theta(t, t))/\langle K \rangle$.

Governing Equations

434

Our full system of equations becomes

$$S(t) = \mu \sum_{a_u=0}^{\infty} (1-\mu)^{a_u} s(t, a_u)$$

$$s(t, a_u) = \begin{cases} \psi(\Theta(t, a_u)) & a_u < t \\ (1-\rho)\psi(\Theta(t, t)) & a_u \geq t \end{cases}$$

$$I(t) = 1 - S(t)$$

$$\Theta(t, 0) = 1$$

$$\Theta(0, a_u) = 1$$

$$\Theta(t, a_u) = \Theta(t-1, a_u-1) - \tau \Phi_I(t-1, a_u-1) \qquad t, a_u \geq 1$$

$$\Phi_I(t, a_u) = \Theta(t, a_u) - \Phi_S(t, a_u)$$

$$\Phi_S(t, a_u) = (1-p_b)^{a_u} \phi_S(t, a_u, a_u) + p_b \sum_{a_e=0}^{a_u-1} (1-p_b)^{a_e} \phi_S(t, a_u, a_e)$$

$$\phi_S(t, a_u, a_e) = \Theta(t - a_e, a_u - a_e)\chi(t, a_e)$$

$$\chi(t, a_e) = \begin{cases} (1-\rho)\frac{\psi'(\Theta(t,t))}{\langle K \rangle} & a_e \geq t \\ \\ (1-P_e)\frac{\psi'(\Theta(t,a_e))}{\langle K \rangle} \\ + P_e \mu \sum_{A_v=a_e+1}^{t-1} (1-\mu)^{A_v-a_e-1} \Theta(t-a_e, A_v-a_e)\frac{\psi'(\Theta(t,A_v))}{\langle K \rangle} & a_e < t \\ + P_e(1-\rho)\Theta(t-a_e, t-a_e)\frac{\psi'(\Theta(t,t))}{\langle K \rangle}(1-\mu)^{t-a_e-1} \end{cases}$$

We can derive a differential equations version of this by replacing the time step n with the time step $n\Delta t$, and then calculating derivatives in the $\Delta t \to 0$ limit.

A differential equations version of this can be derived by treating the time step as Δt rather than 1 and assuming that the event probabilities are all proportional to Δt. Then taking $\Delta t \to 0$ yields differential equations. We will use $\hat{\mu} = \lim_{\Delta t \to 0} \mu/\Delta t$ and similarly define other variables.

435
436
437
438
439
440

We find

$$S(t) = \mu e^{-a_u \mu} s(t, a_u)$$

$$s(t, a_u) = \begin{cases} \psi(\Theta(t, a_u)) & a_u < t \\ (1-\rho)\psi(\Theta(t, t)) & a_u \geq t \end{cases}$$

$$I(t) = 1 - S(t)$$

$$\Theta(t, 0) = 1$$

$$\Theta(0, a_u) = 1$$

$$\left(\frac{\partial}{\partial t} + \frac{\partial}{\partial a} \right) \Theta(t, a_u) = -\tau \Phi_I(t, a_u)$$

$$\Phi_I(t, a_u) = \Theta(t, a_u) - \Phi_S(t, a_u)$$

$$\Phi_S(t, a_u) = e^{-p_b a_u} \phi_S(t, a_u, a_u) + p_b \int_0^{a_u} e^{-p_b a_v} \phi_S(t, a_u, a_e) \, \mathrm{d}a_e$$

$$\phi_S(t, a_u, a_e) = \Theta(t - a_e, a_u - a_e) \chi(t, a_e)$$

$$\chi(t, a_e) = \begin{cases} (1-\rho)\frac{\psi'(\Theta(t,t))}{\langle K \rangle} & a_e \geq t \\[2ex] (1 - P_e)\frac{\psi'(\Theta(t,a_e))}{\langle K \rangle} & \\ + P_e \mu \int_{a_e}^t e^{-\mu(A_v - a_e)} \Theta(t - a_e, A_v - a_e) \frac{\psi'(\Theta(t,A_v))}{\langle K \rangle} \, \mathrm{d}A_v & a_e < t \\ + P_e(1-\rho)\Theta(t - a_e, t - a_e)\frac{\psi'(\Theta(t,t))}{\langle K \rangle} e^{-\mu(t - a_e)} & \end{cases}$$

The simplest numerical method to solve this system of equations would apply an Euler method, which corresponds to solving the discrete-time equations above.

441

442

References

1. Murray CJ, Ortblad KF, Guinovart C, Lim SS, Wolock TM, Roberts DA, et al. Global, regional, and national incidence and mortality for HIV, tuberculosis, and malaria during 1990–2013: a systematic analysis for the Global Burden of Disease Study 2013. The Lancet. 2014;384(9947):1005–1070.

2. Morris M, Kretzschmar M. Concurrent partnerships and transmission dynamics in networks. Social Networks. 1995;17(3):299–318.

3. Morris M, Kretzschmar M. Concurrent partnerships and the spread of HIV. AIDS. 1997;11(5):641–648.

4. Mah TL, Halperin DT. Concurrent sexual partnerships and the HIV epidemics in Africa: evidence to move forward. AIDS and Behavior. 2010;14(1):11–16.

5. Lurie MN, Rosenthal S. Concurrent partnerships as a driver of the HIV epidemic in sub-Saharan Africa? The evidence is limited. AIDS and Behavior. 2010;14(1):17–24.

6. Mah TL, Halperin DT. The evidence for the role of concurrent partnerships in Africa's HIV epidemics: a response to Lurie and Rosenthal. AIDS and Behavior. 2010;14(1):25–28.

7. Morris M. Barking up the wrong evidence tree. Comment on Lurie & Rosenthal, "Concurrent partnerships as a driver of the HIV epidemic in sub-Saharan Africa? The evidence is limited". AIDS and Behavior. 2010;14(1):31–33.

8. Epstein H. The mathematics of concurrent partnerships and HIV: A commentary on Lurie and Rosenthal, 2009. AIDS and Behavior. 2010;14(1):29–30.

9. Lurie MN, Rosenthal S. The concurrency hypothesis in sub-Saharan Africa: Convincing empirical evidence is still lacking. Response to Mah and Halperin, Epstein, and Morris. AIDS and Behavior. 2010;14(1):34–37.

10. Kretzschmar M, White RG, Caraël M. Concurrency is more complex than it seems. AIDS (London, England). 2010;24(2):313.

11. Miller JC, Volz EM. Model Hierarchies in Edge-Based Compartmental Modeling for Infectious Disease Spread. Journal of Mathematical Biology. 2013;67(4):869–899.

12. Epstein H, Morris M. Concurrent partnerships and HIV: an inconvenient truth. Journal of the International AIDS Society. 2011;14(1):13.

13. Volz EM. SIR dynamics in random networks with heterogeneous connectivity. Journal of Mathematical Biology. 2008;56(3):293–310.

14. Volz EM, Meyers LA. Susceptible–infected–recovered epidemics in dynamic contact networks. Proceedings of the Royal Society B: Biological Sciences. 2007;274(1628):2925–2933.

15. Miller JC. A note on a paper by Erik Volz: SIR dynamics in random networks. Journal of Mathematical Biology. 2011;62(3):349–358.

16. Miller JC, Slim AC, Volz EM. Edge-Based Compartmental Modelling for Infectious Disease Spread. Journal of the Royal Society Interface. 2012;9(70):890–906.

17. Miller JC, Volz EM. Incorporating Disease and Population Structure into Models of SIR Disease in Contact Networks. PLoS ONE. 2013;8(8):e69162.

18. Leung K, Kretzschmar M, Diekmann O. Dynamic concurrent partnership networks incorporating demography. Theoretical population biology. 2012;82(3):229–239.

19. Leung KY, Kretzschmar M, Diekmann O. {SI} infection on a dynamic partnership network: characterization of R_0. Journal of mathematical biology. 2014;71(1):1–56.

20. Pinkerton SD. Probability of HIV transmission during acute infection in Rakai, Uganda. AIDS and Behavior. 2008;12(5):677–684.

21. Pilcher CD, Joaki G, Hoffman IF, Martinson FE, Mapanje C, Stewart PW, et al. Amplified transmission of HIV-1: comparison of HIV-1 concentrations in semen and blood during acute and chronic infection. AIDS (London, England). 2007;21(13):1723.

22. Liljeros F, Edling CR, Amaral LAN, Stanley HE, Åberg Y. The web of human sexual contacts. Nature. 2001;411(6840):907–908.

23. Karrer B, Newman MEJ. Message passing approach for general epidemic models. Physical Review E. 2010;82:016101.