Published in final edited form as:

AIDS Behav. 2012 February ; 16(2): 312–322. doi:10.1007/s10461-010-9858-x.

Concurrent Partnerships, Acute Infection and HIV Epidemic Dynamics Among Young Adults in Zimbabwe

Steven M. Goodreau,

Department of Anthropology, University of Washington, Box 353100, Seattle, WA, USA

Susan Cassels,

Department of Epidemiology, University of Washington, Seattle, WA, USA

Danuta Kasprzyk,

Battelle Centers for Public Health Research and Evaluation, Seattle, WA, USA

Daniel E. Montaña,

Battelle Centers for Public Health Research and Evaluation, Seattle, WA, USA

April Greek, and

Battelle Centers for Public Health Research and Evaluation, Seattle, WA, USA

Martina Morris

Departments of Statistics and Sociology, University of Washington, Seattle, WA, USA

Abstract

This paper explores the roles of acute infection and concurrent partnerships in HIV transmission dynamics among young adults in Zimbabwe using realistic representations of the partnership network and all published estimates of stage-specific infectivity. We use dynamic exponential random graph models to estimate partnership network parameters from an empirical study of sexual behavior and drive a stochastic simulation of HIV transmission through this dynamic network. Our simulated networks match observed frequencies and durations of short- and long-term partnerships, with concurrency patterns specific to gender and partnership type. Our findings suggest that, at current behavior levels, the epidemic cannot be sustained in this population without both concurrency and acute infection; removing either brings transmission below the threshold for persistence. With both present, we estimate 20–25% of transmissions stem from acute-stage infections, 30–50% from chronic-stage, and 30–45% from AIDS-stage. The impact of acute infection is strongly moderated by concurrency. Reducing this impact by reducing concurrency could potentially end the current HIV epidemic in Zimbabwe.

Keywords

HIV/AIDS; Sexual networks; Concurrent partnerships; Acute infection; Zimbabwe; Mathematical modeling; Stochastic simulation; ERGM

Introduction

HIV infection is defined by three general stages of disease—acute infection, chronic infection, and AIDS—coinciding with variation in viral load, CD4 cell count, and clinical manifestations [1–4]. It has been speculated for many years that a large fraction of transmissions occur during the short window of acute infection [5–8]. More recently, the first (and to date, only) solid empirical estimates of stage-specific per-coital act transmission probabilities were published [9], based on a landmark longitudinal study of discordant couples in Uganda. Different interpretations of the data have led to a number of different estimates of both stage-specific infectivity [10–12], and the fraction of new HIV infections attributable to the acute stage, ranging from 1 to 41% of infections in different models and populations [10–19] (Table 1). Another approach based on phylogenetics, largely among homosexually active men, suggests that nearly half of transmissions stem from acute-phase cases [20–22].

For more than a decade, some researchers have also pointed out that there is a potentially important interaction between concurrency and acute infection in HIV transmission dynamics [12,23,24]. This is because concurrency ensures that someone newly infected with HIV has the opportunity to transmit to another person during that brief window of peak infection. Very high rates of partner change serve a similar function in core group transmission dynamics. But high rates of partner change are not observed in the general heterosexual population, which makes the synergy between concurrency and acute infection a compelling hypothesis for explaining generalized heterosexual epidemics.

A recently published paper in this journal [14] examined the interaction between concurrency and acute infection, using a modified version of the original 1997 network model used by Morris and Kretzschmar [24]. The authors modified the original model by adding stage-specific transmission probabilities, mortality, and a sensitivity analysis based on the point prevalence of concurrency observed in the Manicaland cohort study in 1998–2000 [25]. They find a strong interaction effect: the impact of concurrency is much greater in the presence of acute infection, and the fraction of infections due to the acute stage is much higher in the presence of concurrency. This paper also demonstrates that the effects of concurrency reported in the original Morris and Kretzschmar paper are not simply artifacts of unrealistic assumptions, as claimed by a recent paper [26]. In fact, as one moves from the early “proof of concept” models to models that are more realistically parameterized, the effects of concurrency become larger, not smaller.

In this paper we extend these recent findings—using a more detailed and empirically based model of the underlying sexual network and observed rates of coital frequency, and we examine the robustness of the findings to measurement error in both the sexual behavior and the stage-specific infectivity estimate.

Methods

The sexual behavior data we use come from the 2005 pre-intervention baseline survey of young adults in ZiCHIRE, the Zimbabwe arm of a five-nation “Popular Opinion Leader” behavioral intervention study. The overall study and Zimbabwe-specific additions have been described elsewhere [27,28]. Briefly, respondents were recruited in thirty rural government-designated “growth points” around Zimbabwe. Inclusion criteria were: being 18–30 years old, having lived in the community at least 2 years, and not being away from the community more than a total of 3 months out of the year. Our analysis was limited to those reporting ever having sex, who were the only ones included in all relevant questions (see the Discussion for implications). Recruitment occurred through local social venues (e.g. stores)

identified during earlier ethnographic work. Approximately three-quarters of respondents were Shona, and one-quarter Ndebele. Mean age was 23.8.

The Zimbabwe study collected both standard sexual behavior summary data and egocentrically sampled network data, in an interviewer-administered questionnaire. Respondents reported the number of sexual partners they had had in the previous 12 months by relationship type. This allows us to distinguish cohabiting partners from other types of partners in the model. For up to the four most recent partners per category, respondents were asked for date (by month) of the first and the last time they had sex with this person. These are the questions recommended by the UNAIDS reference group on Estimates, Modeling and Projection for measuring concurrency [29]. The size of the sample that completed all necessary items for this questionnaire was 4,198 (2,027 females and 2,171 males). We use these data to tabulate the number of partners people have at a given point in time, by sex and partnership type, and to estimate partnership duration by type. These summary statistics are used in the first phase of analysis to estimate the network model, and they are stochastically matched in the network simulation phase.

Point-Partnership Distributions

A person's "point-partnership count" is their number of ongoing partnerships at a given point in time. The data allow us to estimate these counts at the start of each month in the last year. Since most sources of bias in self-reported sexual behavior data would lead to underestimating the true prevalence of concurrency, we use the month with maximum point prevalence for our estimate: 2 months before interview. The number of active partnerships reported by women and men at month two is consistent with estimates of the population sex ratio [30], suggesting the data have good internal validity. However, women classify a higher fraction of partnerships as cohabiting (54.9%) than do men (31.0%), a pattern also observed in a Tanzanian study [31]. We adjust the counts by taking the mean of the two reports; more detail on this step is available in the Supplemental Digital Content. The resulting distribution of partnerships by sex, point partnership counts, and relationship types is shown in Table 2. Note that each cell entry represents the proportion of the *total* population of that sex that has *c* cohabiting partners and *n* non-cohabiting partners in the cross-section, for a given combination of *c* and *n*. We then scale the data in Table 2 to construct a network of 2,000 men and 2,000 women in which those same proportions hold, using a simulated annealing algorithm for relational data [32].

Partnership Durations

Naïve Kaplan–Meier estimates of the average durations for the partnerships active at the start of month two were 45.3 months (for cohabiting) and 11.2 months (otherwise). Given the bias towards observing longer partnerships at a cross-sectional point in time, we re-weighted partnerships according to their relative inclusion probability, which is equivalent to using the harmonic mean of durations. This gives expected average durations of 16.4 and 2.2 months for cohabiting and non-cohabiting partnerships respectively. More sophisticated statistical adjustment methods show comparable dramatic declines in duration estimates once the length-biased sampling effect is removed [33]. In addition, the sample (and model population) is aged 18–30, with mean age 23.8, which also contributes to short estimated partnership durations.

Network Model Estimation

The above statistics are used to estimate the generative parameters of a network, using the exponential random graph model (ERGM) framework [34–41] with an extension for dynamic networks [42]. Dynamic ERGMs are used to predict the probability of a partnership forming or dissolving. The model form looks similar to logistic regression, and

in certain cases simplifies to this, but can also handle complex forms of recursive dependence among observations. That dependence is the hallmark of network models. We specify the partnership formation part of the ERGM with two individual attributes (sex, the point-partnership count, and their interaction) and one relational attribute (cohabiting or not). Represented in conditional auto-logistic form the model is:

$$\begin{aligned} \text{logit} \left(p \left(y_{ij}=1 | Y^c \right) \right) = & \theta \delta(e) + \sum_s \theta_{s,1,1} \delta(u_{s,1,1}) \\ & + \sum_s \theta_{s,1} \delta(u_{s,1}) + \theta_{m,2} \delta(u_{m,2}) \\ & + \theta_{f,2+} \delta(v_{f,2+}) + \theta_{m,3+} \delta(v_{m,3+}) \end{aligned}$$

where y_{ij} = the pair of persons i and j , and $y_{ij} = 1$ indicates they are partners; Y^c = the rest of the pairs in the network, excluding the y_{ij} pair; e = total number of partnerships of all types in the network; $u_{s,n,r}$ = # of persons of sex s with exactly n partnerships of type r (1 = cohabiting, 2 = non-cohabiting); $u_{s,n}$ = # of persons of sex s with exactly n total partnerships of all types; $v_{s,n+}$ = # of persons of sex s with n total partnerships of all types.

The function δ represents the change in the specific network statistic (e , u or v) when the i,j pair become partners.

The terms in the partnership formation model thus fit the following summary statistics: the total number of partnerships, the number of men and women with exactly one cohabiting partner and the number with exactly one partner in total, the number of men with exactly two total partners and with three or more total partners, and the number of women with two or more total partners. The model for pair dissolution is a mixed Bernoulli, with separate terms for each partnership type, based on the estimates of duration. This generates an exponential survival curve for each type of partnership. Model constraints include $v_{f,2,1} = v_{m,3,1} = 0$; that is, no women have more than one cohabiting partner, and no men have more than two (see Supplemental Digital Content for more detail on the model specification). Together, these model terms and constraints provide a detailed specification of concurrency patterns by sex and partnership type.

Partnership Network Simulation

In the ERGM framework, the network simulation is driven by the same model as the one used for estimation. This ensures that the expected values of the network statistics in the model (which equal the values observed in the data) are preserved at every cross-sectional time point, while the values themselves vary stochastically around these expectations.

We implement network simulation and estimation using *statnet* (<http://www.statnetproject.org>) [43]. Fitting details used (e.g. burn-in, MCMC sample constraints) are available from the authors on request. Simulated dynamic networks were checked to ensure that they retained the cross-sectional structure and relational durations derived from the data.

Epidemic Simulation

Transmission is modeled in monthly time steps. For each time step, the probability of transmission in discordant couples is determined by the infection stage of the positive partner and coital frequency. Given current uncertainty about stage-specific transmission probabilities, we treat this as a sensitivity parameter. We consider four sets of transmission probabilities from commonly cited articles in the literature on heterosexual HIV transmission dynamics: “Wawer” [9], “Pinkerton” [10], “Abu-Raddad” [11], and

“Hollingsworth” [12]. All of these estimates derive from the same data [9], but they define and estimate the number of stages, the duration of stages, and the stage-specific transmission probabilities from the data in different ways. For instance, Hollingsworth and Pinkerton both have a single acute stage of around 3 months, Abu-Raddad a single 2.5-month acute stage, and Wawer a 5-month acute stage followed by an additional 10-month stage of intermediate transmissibility. Full details for each model can be found in the Supplemental Digital Content. The Hollingsworth scenario is the most parsimonious; it uses three stages and assumes no coital acts in the final 10 months of AIDS. The other three scenarios use seven or eight stages (with varying definitions), and make stage-specific assumptions about coital frequency that generate monthly transmission probabilities per partnership. The original source papers for these scenarios do not report the number of coital acts in the final 5 months, so we treat this as another sensitivity parameter, with three levels for each scenario (excluding Hollingsworth): High = 6.2/month, Middle = 3.1/month, Low = 0, based on reported coital acts during the 5 months prior to the final 5 months. We assume mean survival of 122 months after infection [12]. Current estimates suggest the population growth rate in Zimbabwe is close to zero [44], so we model a stable population here; deaths are immediately replaced with new (uninfected) arrivals of the same sex.

We start the epidemic with one male and one female infected, and run until equilibrium is reached for each scenario. The alternative, beginning with a level of HIV prevalence similar to the present-day, would require us to assume a distribution for the time since infection for the initially infected population, and we have little information on this for a realistic population. Starting with two seeds means that we are simulating the epidemics that would emerge at equilibrium under the behavioral conditions observed in 2005. There is ample evidence that sexual behavior has changed in Zimbabwe since the epidemic began [45], and in many other African countries with generalized epidemics [46], so the initial transitory dynamics in these simulations will not be (and are not meant to be) an accurate representation of the original epidemic trajectories. What these models do show is the equilibrium behavior of the epidemic under the sexual behavior conditions observed in 2005.

The equilibrium prevalence from simulations started with two seeds would be no different than those started at a higher prevalence (if the stage-specific infection distribution in the initially infected population could be properly specified). In this type of stationary epidemic model, there are direct relationships among successful invasion, persistence and the endemic prevalence at equilibrium. This is traditionally captured by the summary statistic that represents the initial reproductive ratio of infection, R_0 . Our model, with vital dynamics, a stable population size and no change in the underlying behavioral or biomedical transmission parameters, is structurally an SI model with birth and death. Epidemic theory shows that epidemic prevalence for the homogenous version of this model when $R_0 > 1$ is approximated by $1 - 1/R_0$ [47]. With complex network structures, we no longer have an analytic solution for R_0 . However, a fundamental insight of modeling theory is that for any epidemic without change in behavioral or biomedical transmission parameters, the reproductive ratio at any timepoint with disease prevalence > 0 is less than it is at introduction to an entirely susceptible population [48]. The consistent inability of an epidemic to invade in a population implies that $R_0 < 1$, which implies that the reproductive ratio for any level of prevalence is also less than one. A scenario that produces runs that are consistently unable to invade would therefore also be unable to sustain an epidemic, regardless of the initial prevalence at which it is started.

Epidemic simulations based on deterministic models produce a single equilibrium prevalence value—either disease-free (when the basic reproductive ratio $R_0 < 1$) or endemic (when $R_0 > 1$). This makes it easy to classify whether a given scenario is above or below the

reproductive threshold. With stochastic models, the equilibrium prevalence is not represented by a single value, but by a distribution of prevalence values across many realizations. Stochastic models also introduce the possibility that a scenario with parameters that support a positive endemic equilibrium will still generate some individual runs that go to extinction; this is particularly common when initial prevalence is small, as it is here. Thus the classification of a scenario as above or below the reproductive threshold is based on observing whether any runs persist. Scenarios in which some percentage of runs successfully invade and then sustain at endemic levels for centuries clearly possess an endemic equilibrium. For those scenarios in which infection never invades or sustains across a single run out of 1,000, we infer that $R_0 < 1$ and there is no epidemic potential. Jacquez and Simon [47] provide an analytical and simulated estimate of $(1/R_0)^k$ for the expected probability of non-invasion for the stochastic homogeneous SI model with birth and death, where k is the number of seeds. This value would imply, for example, that a scenario with two initial seeds and a mean endemic prevalence of 10% (conditional on non-extinction) should see about 81% of simulations go extinct. For our version with network structure, we again have no exact analytical solutions; nevertheless, for our models in which some runs sustain, the extinction probability approximates the simpler analytical solution well. The behavior we observe for the remainder of our scenarios (no cases of invasion after 1,000 simulations with two seeds in a population of size 2,000) is qualitatively different than this, and clearly indicates that $R_0 < 1$.

We therefore report two sets of epidemic outcomes: the fraction of runs that successfully invade, and the average equilibrium prevalence for runs that persist. The scenarios we investigate are summarized below:

- *Baseline scenario:* For this we use the Hollingsworth estimates—the most parsimonious and specific of the infectivity estimates—and the observed network statistics. We refer to this as the “baseline” scenario because we will use it as the basis for comparison for all subsequent models.
- *Rounded-up scenario:* To test the sensitivity of this model to potential under-reporting of sexual partnerships, we conduct a set of runs rounding up the cross-sectional mean number of partners per person to one significant digit, from 0.66 to 0.7, and compare to the baseline.
- *Stage-specific infectivity variants:* These include the Pinkerton, Wawer and Abu-Raddad scenarios, each with three versions for coital frequency in the last 5 months.
- *Counterfactual scenarios:* We consider two counterfactual (“what-if”) scenarios, each using the Hollingsworth transmission estimates:
 - *No concurrency:* This examines the role that concurrency plays in amplifying the impact of acute-stage transmission by eliminating concurrent partnerships. For this scenario, we estimated an ERGM for a network with the same total number of partnerships and the same durations as observed in the Zimbabwe data, but constrained the point-partnership counts to 0 or 1. Networks simulated from this model preserve the same number of partnerships of each type (at any time point and over time), the same total time spent in partnerships, the same relational durations, and the same coital frequency within partnerships as the baseline run, so the comparison isolates the effect of concurrency alone.
 - *No acute peak:* This examines the potential impact of reducing the acute-stage transmission probability when there is no behavior change. For this scenario, we modify the Hollingsworth parameters by reducing the acute-stage

probability (0.21% per serodiscordant partnership per month) to the chronic-stage level (0.0088%).

Results

The prevalence of concurrency in this population is relatively high, compared to levels observed in the United States [49,50]. Overall, the point prevalence is 7.3% and the cumulative annual prevalence is 12.3%, with a clear gender gap: 2.8% vs. 11.6% for point prevalence among women and men respectively, and 4.4% vs. 19.7% for cumulative annual prevalence. Women were more likely than men to report exactly one ongoing partnership at a point in time (55.6% vs. 43.4%), while men reported higher rates of both >1 ongoing partnership (11.6% vs. 2.8%), and no ongoing partnerships (45.0% vs. 41.6%).

Networks consistent with these distributions are quite sparse in the cross-section; most members will reside in components of 1–2 people. Figure 1 displays a random cross-sectional realization drawn from the distribution of networks specified by an ERGM fitting these statistics to 2,000 persons. Cross-sectional connectivity, however, does not exclusively determine the reachable path of infection in a dynamic network [49,51]. Epidemic invasion and persistence also depends on the forward path from each initial infected case, which is defined by partnership durations and sequence. Still, the sparseness in the cross-sectional network does suggest that epidemics may not be very robust in this population. We can evaluate this from the stochastic variability in epidemic persistence across repeated runs.

Summaries of the epidemic outcomes for all scenarios are shown in Table 3. The baseline parameter set—using the Hollingsworth stage-specific infectivity estimates and the observed network statistics—will sustain an epidemic in this population, with average equilibrium prevalence of 8.8% (range: 1.3–12.1%). Epidemic invasion was not very robust, however; most runs on these networks died out early on. On average the baseline networks needed 11 runs before the first sustained epidemic was produced (though this too was variable, with a range of 1–37). All of this, in turn, was quite sensitive to a small increase in number of partnerships active on each day. With the “rounded-up” scenario (which increased the number of partnerships by 6% and concurrency by 2%), average prevalence rose to 13.8%, was much less variable, with a range of 12.3–14.6%, and an average of only 5 runs was needed before the first persistent epidemic. Together, this suggests the baseline scenario is very close to the threshold for epidemic invasion and persistence.

The other stage-specific infectivity variants produced smaller epidemics, with more variable outcomes, and were extremely sensitive to the assumptions about coital frequency during the AIDS stage. The Pinkerton scenario was the only one to produce any sustained epidemics under the low frequency assumption, but very few runs persisted—it took an average of 98 runs before the first sustained epidemic was observed—and equilibrium prevalence averaged only 1% when sustained. The Abu-Raddad scenario failed to produce a single persistent epidemic after 1,000 runs (100 runs each on 10 network realizations) for both the medium and low frequency assumptions. Recall that the baseline Hollingsworth scenario assumes the lowest value (no transmission activity), for twice as long (10 months).

Since the results from these variants are not very robust, we focus primarily on the findings from the baseline model for the remaining comparisons.

The proportion of transmissions due to each stage over time in the baseline scenario is shown in Fig. 2. The first panel (Fig. 2a) shows that these fractions stabilize quickly. On average, the stage-specific fractions are 22%, 49%, and 29% from acute-, chronic- and AIDS-stage infections, respectively, with remarkably little variation across runs. The second panel (Fig. 2b) shows the cumulative fractions of the age of all source infections in the ten

sustained runs (note the x-axis is not time, but the age of the source infection). This, too, shows little variation. By comparison, the Pinkerton and Wawer scenarios produce similar acute-stage fraction estimates, the Abu-Raddad runs produce a lower estimate, and all three variants predict more infections from the AIDS stage. More detailed results for each model are in the Supplemental Digital Content.

Counterfactual Scenarios

The no-concurrency scenario fails to create a single persistent epidemic in all 1,000 runs. That is, if the same number of partnerships reported in the survey were to occur sequentially instead of concurrently, with the same durations and levels of coital frequency, the epidemic would not be sustainable in this population at 2005 behavior levels.

The no-acute-peak scenario also fails to generate a persistent epidemic in any of 1,000 runs. Thus, even though the acute stage only contributes about 20% of infections in sustained epidemics (across the various infectivity models), reducing transmissibility during acute infection down to chronic levels makes the epidemic unsustainable in this population at the 2005 behavior levels.

If we restrict the focus to the first 5% of the time series for the counterfactual runs that survive this long, the fraction of infections due to the acute phase drops from 22% to 11% in the no-concurrency scenario, and to 0% in the no-acute-peak scenario.

Discussion

Our findings suggest that the epidemic in Zimbabwe is very close to the persistence threshold—small changes in either behavior or infectivity may be enough to push it into eventual extinction. Differences as small as the interpretation of stage-specific infectivity used by researchers working with the same underlying source data were sufficient to push the transmission dynamics across that threshold, while a slight increase of 6% in the number of active partnerships might be enough to ensure persistence at prevalence >10%. Both of these suggest that measurement error, in both biological and behavioral parameters, may play a significant role in the projections from epidemic models in Zimbabwe. This uncertainty should be acknowledged in simulation studies; much more attention should be paid to both the validity and reliability of the biological and behavioral measures, and to the accurate representation of all elements of the transmission system.

The one finding for which there is very little uncertainty comes from our counterfactual scenarios: both concurrency and the peak in acute infection are needed for epidemic persistence of HIV in this population at this point. Neither alone is sufficient. While some of the other scenarios only generate one persistent epidemic for every 10 or 100 tries, both counterfactuals failed to generate a single epidemic in 1,000 tries, which strongly suggests these scenarios lack epidemic potential. This is a qualitative difference. Only a small fraction of infections directly stemmed from acute index cases, but eliminating this peak infection window makes the epidemic unsustainable. The same is true for concurrent partnerships; eliminating concurrency, while keeping the same number of total partners, partnership durations, and coital frequencies, and the same peak infection window, leads to epidemic extinction. The joint impacts of concurrency and acute infection, long the subject of speculation, are clearly confirmed here.

Our estimates of the fraction of infections associated with the acute-stage do not always match the original published estimates. Hollingsworth et al. [12] estimate 31%, and Abu-Raddad and Longini [11] estimate ~5–13% of transmissions stem from acute-stage infections at equilibrium. Using their infectivity profiles, we estimate about 20% for the

Hollingsworth scenario and 11% for the Abu-Raddad scenario. Pinkerton [10] estimates that 89.1% of the transmissions during the first 20 months of infection occur in the acute phase; our model estimates 76% using Pinkerton's parameters. Since we are using the same infectivity parameters as these studies in each comparison, the differences in our estimates must stem from other components of our models.

With respect to these papers, the critical difference is the way we model behavior. Our behavioral model is different in two important ways. First, our transmission network parameters are estimated from data on a single population. This is not the case for the models used in other studies. Parameters in those models are taken from a number of different populations, in order to meet the input needs of the simulation setup given the constraints of empirical datasets. There are real differences in behavior within and between the populations of sub-Saharan Africa that make this mixed input approach hard to defend. Restricting behavioral inputs to a single population does limit generalization, especially to "sub-Saharan Africa." But in return, it strongly increases the validity of inferences to the population of interest. Second, our model accurately represents the observed timing and sequence of multiple partnerships (both concurrent and sequential), and the distinctions between types of partners. This is also not the case in the other models. The recent paper by Eaton et al. [14] that included a stylized version of concurrency and used the Hollingsworth infectivity estimates, by contrast, show similar results to ours. They find that 16–28% of infections are attributable to the acute phase at equilibrium across a range of concurrency levels. Both our study and theirs examine a "no concurrency" counterfactual scenario, and both find that it makes a qualitative difference in the epidemic outcome: in this population, it is not possible to sustain an epidemic without concurrency, if the empirical distribution of the number and types of partnerships is represented with fidelity.

This has an important implication for using models to understand the epidemiology of HIV in Africa: details matter, but some details matter more than others. In order to match observed epidemic trajectories, previous modeling studies that have not explicitly represented concurrency have had to posit absurdly high rates of partner accumulation. One example from a widely cited study has to assume that the *average* male and female in the population of Yaoundé, Cameroon has 221 lifetime partners, and 7% have an average of 2,870 partners, each relationship at least 1 week long, over 35 sexually-active years [11]. In another widely cited study, the average person in "sub-Saharan Africa" (male and female both) is assumed to have 120 partners over their lifetime, and the top 10% of both sexes have 723 partners over 40 years, each with 25 sex acts or more [52]. These are not extreme examples of model assumptions; they are typical for models that do not explicitly represent concurrency. Behavioral patterns like this are not found in general heterosexual populations anywhere, but they are apparently necessary, in the absence of concurrency, to generate realistic epidemics. One reason is that the mathematics underlying these examples implicitly assumes that a partnership can only transmit if it is serodiscordant at its outset, and not if it becomes serodiscordant during its course. The profound disconnect between these types of assumptions and empirical data has been criticized [53], and has been used to argue for the importance of non-sexual transmission routes in Africa [54]. Our modeling approach suggests a different interpretation. Consistent with other emerging studies [14,24], we find that far fewer partners are necessary to generate substantial HIV epidemics—when concurrency is explicitly measured and modeled. While this provides strong support for the role of sexual transmission of HIV, we suspect that the massive exaggeration of the number of partners in these other models is bound to lead to artifacts in their primary findings. This suggests it would be worth going back to re-evaluate the validity of the findings from key modeling studies in the past decade.

Our model also has many assumptions and limitations that must be kept in mind when making inferences. Our model excluded people who had never had sex (19.9% of the original study sample), as well as commercial sexual contacts. Each would affect our equilibrium prevalence estimates, though in opposite directions. Including the not-yet-sexually-active population would lower HIV prevalence estimates somewhat; if we chose to model them as a distinct, permanent class, as is sometime done, then the new estimates would simply be 80.1% of the values previously estimated. Including commercial contacts would, of course, raise prevalence estimates. Our data were drawn from young adults, and inference should be focused on this subpopulation. We did not consider additional network structuring induced by geographic mixing, nor mixing by age or other exogenous demographic or social variables. Our modeling framework can represent these, but our goal here was more limited. We assumed that deaths were balanced with new arrivals due to limitations on the ERGM framework that were in place when we began this research; these limitations are now lifted [42], enabling future work to explore the interactions of vital dynamics and relational structure in more detail. We did not consider variation in circumcision status or co-infection with STDs or other infections. Neither STD treatment nor circumcision was needed to push our epidemic into extinction in this study, but their absence may help to keep transmission marginally above the reproductive threshold in the population. Our empirical estimates of concurrency are subject to several sources of measurement error. Some may lead to overestimates (e.g., our selecting the month with highest prevalence), others to underestimates (e.g. social desirability bias in self-reported sexual behavior data). Our sensitivity analyses suggest these could well be important for accurate inference. Our estimates of stage-specific transmission probabilities and coital frequency are based on data from a single study, and are also subject to measurement error (this limitation applies to all other published studies on this topic as well). There is uncertainty in the duration and magnitude of the acute infection window, and in coital frequency during late-stage AIDS. For these limitations, the sensitivity analyses we conduct on transmission probabilities gives some intuition regarding the impact on model outcomes.

Finally, our models do not accurately reproduce the rapid rise in prevalence that was observed in Zimbabwe, from the low single digits in the early 1980s to around 20% in the early 2000s; they took well over a century to achieve equilibrium prevalence when starting from two initially infected cases. This would be a serious limitation if we were trying to replicate the original epidemic trajectory. However, that was not our goal. Rather, time in our model simply represents a “burn-in” period required to determine the equilibrium conditions implied if the behavior observed in our data had and always will be present. Reproducing the original trajectory would require an accurate model of behavior in the pre-1980 period, for which we have no data at all, let alone the type of egocentric network data needed to accurately represent concurrency. There is evidence from the longitudinal cohort study in Manicaland that risky behavior declined from 1998 to 2003, with reductions in both casual partners and concurrency, especially among men [45]. But this is nearly two decades after the start of the epidemic. It is worth noting that Eaton et al., using the 1998–2000 Manicaland data, also find that while the final prevalence in their simulated epidemics is in a realistic range (9–23%, depending on the level of concurrency), it takes 50–100 years to reach these levels, even though they start with a much higher fraction of infections (1%, 100 males and 100 females). They draw the conclusion that something other than concurrency must have caused the early explosive epidemic growth. Since their models use behavioral data collected 20 years after the epidemic began, this inference cannot really be drawn. Unfortunately, we will never have the data we need to answer this question empirically. What we can do now, in the absence of data, is to use this modeling framework to identify the sexual behavior conditions under which a major epidemic *could* unfold in one or two decades. This is an important topic for future research.

Despite these limitations, we believe our findings have important implications for HIV prevention. Logic dictates that a short acute infection window depends on either rapid serial partner acquisition or concurrency to generate epidemic persistence. These patterns of relational timing vary widely among populations in which HIV is spread, but concurrency is the only plausible driver in generalized epidemics. Of the two necessary components for persistence that we found here—peak infection and concurrency—concurrency is something we can address now. And it may be all we need to address, especially in populations that are, as Zimbabwe seems to be, close to the threshold for epidemic persistence. A vaccine only needs to achieve the critical coverage level for herd immunity to kick in, and in this sense, concurrency reduction will function as a behavioral vaccine. We eliminated concurrency in our counterfactual scenario for clarity, but elimination, like 100% vaccine coverage, is not necessary. Using their stylized version of concurrency, Eaton et al. found that rates of concurrency of 4% and below could not sustain an epidemic at the partnership acquisition rates reported in Zimbabwe in 2000. Better models, and better data will help to bracket the target level of concurrency reduction needed to confer herd immunity, but this is a very good start. It is important to recognize that the target level will be population-specific; it depends on the distribution of partnership duration, and gender-specific partnering patterns, so it will be an empirical question for each case. Estimates of this concurrency reduction target are a key component to “knowing your epidemic” and for developing the right prevention package.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors would like to thank the study participants as well as Mark Handcock, David Hunter, Pavel Krivitsky, Carter Butts, and the entire *statnet* development team. We are also grateful for the helpful comments from people who read earlier drafts, including Jim Shelton and Helen Epstein. SMG and SC were supported in part by the Puget Sound Partners for Global Health (Research and Technology Project Award 26145). NIH provided support for the data collection (U10-MH061544, R21-AA014802), the methodology and software development (R01-HD041877 and R01-DA12831), and the analysis (K99-HD057553 and P30-AI27757).

References

1. Coombs RW, Speck CE, Hughes JP, Lee W, Sampoleo R, Ross SO, et al. Association between culturable human immunodeficiency virus type 1 (HIV-1) in semen and HIV-1 RNA levels in semen and blood: Evidence for compartmentalization of HIV-1 between semen and blood. *J Infect Dis.* 1998; 177(2):320–30. [PubMed: 9466517]
2. Leynaert B, Downs AM, de Vincenzi I. Heterosexual transmission of human immunodeficiency virus—variability of infectivity throughout the course of infection. *Am J Epidemiol.* 1998; 148(1): 88–96. [PubMed: 9663408]
3. Little S, McLean A, Spina C, Richman D, Havlir D. Viral dynamics of acute HIV-1 infection. *J Exp Med.* 1999; 190(6):841–50. [PubMed: 10499922]
4. Quinn TC, Wawer MJ, Sewankambo N, Serwadda D, Li CJ, Wabwire-Mangen F, et al. Viral load and heterosexual transmission of human immunodeficiency virus type 1. *N Engl J Med.* 2000; 343(13):921–9. [PubMed: 10738050]
5. Jacquez JA, Koopman J, Simon CP, Longini IMJ. Role of the primary infection in epidemics of HIV-infection in gay cohorts. *J Acquir Immune Defic Syndr Hum Retrovirol.* 1994; 7(11):1169–84.
6. Koopman JS, Jacquez JA, Welch GW, Simon CP, Foxman B, Pollock SM, et al. The role of early HIV infection in the spread of HIV through populations. *J Acquir Immune Defic Syndr Hum Retrovirol.* 1997; 14:249–58. [PubMed: 9117458]
7. Longini IMJ, Clark WS, Haber M, Hosburgh RJ. The stages of HIV infection: waiting times and infection transmission probabilities. *Lect Notes Biomath.* 1989; 83:111–37.

8. Pilcher CD, Tien HC, Eron JJJ, Vernazza PL, Leu S-Y, Stewart PW, et al. Brief but efficient: acute HIV infection and the sexual transmission of HIV. *J Infect Dis.* 2004; 189(10):1785–92. [PubMed: 15122514]
9. Wawer MJ, Gray RH, Sewankambo NK, Serwadda D, Li XB, Laeyendecker O, et al. Rates of HIV-1 transmission per coital act, by stage of HIV-1 infection, in Rakai, Uganda. *J Infect Dis.* 2005; 191(9):1403–9. [PubMed: 15809897]
10. Pinkerton SD. Probability of HIV transmission during acute infection in Rakai, Uganda. *AIDS Behav.* 2008; 12(5):677–84. [PubMed: 18064559]
11. Abu-Raddad LJ, Longini IA. No HIV stage is dominant in driving the HIV epidemic in sub-Saharan Africa. *AIDS.* 2008; 22(9):1055–61. [PubMed: 18520349]
12. Hollingsworth, TD.; Anderson, RM.; Fraser, C., editors. HIV-1 transmission, by stage of infection. Univ. Chicago Press; 2008.
13. Cohen MS, Pilcher CD. Amplified HIV transmission and new approaches to HIV prevention. *J Infect Dis.* 2005; 191(9):1391–3. [PubMed: 15809893]
14. Eaton JW, Hallett TB, Garnett GP. Concurrent sexual partnerships and primary HIV infection: a critical interaction. *AIDS Behav.* 2010 doi:10.1007/s10461-010-9787-8.
15. Hayes R, White R. Amplified HIV transmission during early-stage infection. *J Infect Dis.* 2005; 193:604–5. [PubMed: 16425144]
16. Pinkerton SD. How many sexually-acquired HIV infections in the USA are due to acute-phase HIV transmission? *AIDS.* 2007; 21(12):1625–9. [PubMed: 17630558]
17. Prabhu VS, Hutchinson AB, Farnham PG, Sansom SL. Sexually acquired HIV infections in the United States due to acute-phase HIV transmission: an update. *AIDS.* 2009; 23(13):1792–4. [PubMed: 19684485]
18. Rapatski BL, Suppe F, Yorke JA. HIV epidemics driven by late disease stage transmission. *J Acquir Immune Defic Syndr.* 2005; 38(3):241–53. [PubMed: 15735440]
19. Xiridou M, Geskus R, de Wit J, Coutinho R, Kretzschmar M. Primary HIV infection as source of HIV transmission within steady and casual partnerships among homosexual men. *AIDS.* 2004; 18(9):1311–20. [PubMed: 15362664]
20. Pao D, Fisher M, Hue S, Dean G, Murphy G, Cane PA, et al. Transmission of HIV-1 during primary infection: relationship to sexual risk and sexually transmitted infections. *AIDS.* 2005; 19(1):85–90. [PubMed: 15627037]
21. Yerly S, Vora S, Rizzardi P, Chave JP, Vernazza PL, Flepp M, et al. Acute HIV infection: impact on the spread of HIV and transmission of drug resistance. *AIDS.* 2001; 15(17):2287–92. [PubMed: 11698702]
22. Brenner, BG.; Roger, M.; Routy, JP.; Moisi, D.; Ntemgwa, M.; Matte, C., et al., editors. High rates of forward transmission events after acute/early HIV-1 infection. Univ Chicago Press; 2007.
23. Hudson CP. AIDS in rural Africa: a paradigm for HIV-1 prevention. *Int J STD AIDS.* 1996; 7(4): 236–43. [PubMed: 8876353]
24. Morris M, Kretzschmar M. Concurrent partnerships and the spread of HIV. *AIDS.* 1997; 11(5): 641–83. [PubMed: 9108946]
25. Gregson S, Nyamukapa CA, Garnett GP, Mason PR, Zhuwau T, Carael M, et al. Sexual mixing patterns and sex-differentials in teenage exposure to HIV infection in rural Zimbabwe. *Lancet.* 2002; 359(9321):1896–903. [PubMed: 12057552]
26. Sawers L, Stillwaggon E. Concurrent sexual partnerships do not explain the HIV epidemics in Africa: a systematic review of the evidence. *J Int AIDS Soc.* 2010; 13:34. [PubMed: 20836882]
27. NIMH Collaborative HIV/STD Prevention Trial Group. Methodological overview of a five-country community-level HIV/sexually transmitted disease prevention trial. *AIDS.* 2007; 21(Suppl 2):S3–18.
28. Kasprzyk, D.; Montano, DE. Application of an integrated behavioral model to understand HIV prevention behavior of high-risk men in rural Zimbabwe. In: Ajzen, I.; Hornik, R., editors. Prediction and change of health behavior: applying the reasoned action approach. Lawrence Erlbaum Associates, Inc.; Hillsdale, NJ: 2007.
29. UNAIDS Reference Group on Estimates Modelling and Projections. Consultation on concurrent sexual partnerships: recommendations from a meeting of the UNAIDS reference group on

estimates, modelling and projections held in Nairobi, Kenya, April 20–21 2009.
http://www.epidem.org/Publications/Concurrency%20meeting%20recommendations_Final.pdf2009

30. Central Intelligence Agency. CIA World Factbook—Zimbabwe. [Accessed 12 Dec 2009]. 2009
<https://www.cia.gov/library/publications/the-world-factbook/geos/zi.html>
31. Nnko S, Boerma JT, Urassa M, Mwaluko G, Zaba B. Secretive females or swaggering males? An assessment of the quality of sexual partnership reporting in rural Tanzania. *Soc Sci Med*. 2004; 59(2):299–310. [PubMed: 15110421]
32. Handcock, MS.; Hunter, DR.; Butts, CT.; Goodreau, SM.; Morris, M. ERGM: a package to fit, simulate and diagnose exponential-family models for networks. 2.2-5 ed. 2010.
33. Burington B, Hughes JP, Whittington WL, Stoner B, Garnett G, Aral SO, et al. Estimating duration in partnership studies: issues, methods and examples. *Sex Transm Infect*. 2010; 86(2):84–9. [PubMed: 20332366]
34. Frank O, Strauss D. Markov graphs. *J Am Stat Assoc*. 1986; 81(395):832–42.
35. Strauss D, Ikeda M. Pseudolikelihood estimation for social networks. *J Am Stat Soc*. 1990; 85(409):204–12.
36. Wasserman S, Pattison P. Logit models and logistic regressions for social networks. I. An introduction to Markov graphs and p^* . *Psychometrika*. 1996; 60:401–25.
37. Handcock, MS. Statistical models for social networks: degeneracy and inference. In: Breiger, RL.; Carley, KM.; Pattison, P., editors. *Dynamic social network modeling and analysis: workshop summary and papers*. National Academies Press; Washington: 2003. p. 229–40.
38. Snijders TAB, Pattison PE, Robins GL, Handcock MS. New specifications for exponential random graph models. *Sociol Methodol*. 2006; 36(1):99–153.
39. Robins G, Morris M. Advances in exponential random graph (p^*) models. *Soc Networks*. 2007; 29(2):169–72.
40. Handcock MS, Hunter DR, Butts CT, Goodreau SM, Morris M. Statnet: software tools for the representation, visualization, analysis and simulation of social network data. *J Stat Softw*. 2008; 24(1):1–11. [PubMed: 18612375]
41. Hunter DR, Goodreau SM, Handcock MS. Goodness of fit of social network models. *J Am Stat Assoc*. 2008; 103(481):248–58.
42. Krivitsky, PN. *Statistical models for social network data and processes*. University of Washington; Seattle: 2009.
43. Handcock, MS.; Hunter, DR.; Butts, CT.; Goodreau, SM.; Morris, M. *Statnet: software tools for the statistical modeling of network data*. 2.0 ed. 2003.
44. The World Bank. *World development indicators*. The World Bank; Washington: 2010.
45. Gregson S, Gonese E, Hallett TB, Tarubekera N, Hargrove JW, Lopman B, et al. HIV decline in Zimbabwe due to reductions in risky sex? Evidence from a comprehensive epidemiological review. *Int J Epidemiol*. 2010; 39(5):1311–23. [PubMed: 20406793]
46. Joint United Nations Programme on HIV/AIDS. *AIDS epidemic update*. UNAIDS, World Health Organization; Geneva: 2009.
47. Jacquez JA, Simon CP. The stochastic SI model with recruitment and deaths. 1. Comparison with the closed SIS model. *Math Biosci*. 1993; 117(1–2):77–125. [PubMed: 8400585]
48. May, RM. Population biology of microparasite infections. In: Hallam, TG.; Levin, SA., editors. *Mathematical ecology: an introduction*. Springer-Verlag; Berlin: 1986. p. 405–42.
49. Morris M, Kurth AE, Hamilton DT, Moody J, Wakefield S, The Network Modeling Group. Concurrent partnerships and HIV prevalence disparities by race: linking science and public health practice. *Am J Public Health*. 2009; 99(6):1023–31. [PubMed: 19372508]
50. Morris M, Epstein H, et al. Timing is everything: international variations in historical sexual partnership concurrency and HIV prevalence. *PLoS One*. 2010:e14092. doi:10.1371/journal.pone.0014092. [PubMed: 21124829]
51. Moody J. The importance of relationship timing for diffusion. *Soc Forces*. 2002; 81(1):25–56.

52. Hallett TB, Singh K, Smith JA, White RG, Abu-Raddad LJ, Garnett GP. Understanding the impact of male circumcision interventions on the spread of HIV in southern Africa. *PLoS One*. 2008; 3(5):e2212. [PubMed: 18493593]
53. Deuchert E, Brody S. Plausible and implausible parameters for mathematical modeling of nominal heterosexual HIV transmission. *Ann Epidemiol*. 2007; 17(3):237–44. [PubMed: 17320790]
54. Gisselquist D, Rothenberg R, Potterat J, Drucker E. HIV infections in sub-Saharan Africa not explained by sexual or vertical transmission. *Int J STD AIDS*. 2002; 13:657–66. [PubMed: 12396534]

Fig. 1.

This example network is drawn from the distribution of networks specified by the baseline ERGM. It represents a typical cross-sectional slice from the simulated dynamic network, and preserves the empirical summary statistics observed in the Zimbabwe data

Fig. 2.

a Proportion of new infections by stage over time, averaged across all ten Hollingsworth base model runs. **b** Cumulative density function for transmission events by age of transmitter. The y-axis represents the cumulative number of transmission events from source partners who have been infected no more than the number of months on the x-axis, divided by the total number of transmission events. Each line represents one of the ten Hollingsworth base model runs. The horizontal line at the end of the age axis is due to the assumption of no sex during the last 10 months of infection for this model

Table 1

Previous studies on proportion of HIV infections by stage of index case

Author (year)	Population	% infections from acute stage	Duration of acute stage
Eaton et al. (2010) [14]	Heterosexual, Zimbabwe	16–28%	88 days
Pinkerton (2008) [10]	Heterosexual, Uganda ^a	–	49 days
Hollingsworth et al. (2008) [12]	Heterosexual, Uganda ^a	9–31%	2.9 months
Abu-Raddad and Longini (2008) [11]	Heterosexual, sub-Saharan Africa ^a	13–18% overall 8–15% at equilibrium	2.5 months
Hayes and White (2006) [15]	Heterosexual, Uganda ^a	23–41%	5 months
Cohen and Pilcher (2005) [13]	Heterosexual, Uganda ^a	~50%	5 months
Pinkerton (2007) [16]	US population	8.6% (2.5–17.3%)	49 days
Prabhu (2009) [17]	US population	11.4%	49 days
Xiridou et al. (2004) [19]	MSM, Amsterdam	11%	3 months
Rapatski et al. (2005) [18]	MSM, San Francisco	1.3%	6 months
Brenner et al. (2007) [22]	Majority MSM, Urban Quebec	49.4%	N/A Phylogenetic analysis
Pao et al. (2005) [20]	Majority MSM, United Kingdom	34%	N/A Phylogenetic analysis
Yerly et al. (2001) [21]	Swiss HIV cohort study (42% MSM, 13% IDU)	30%	N/A Phylogenetic analysis

^aInfectivity estimates based on Wawer et al. [9]

Table 2
Adjusted cross-sectional partnership distributions used for network estimation and simulation

(a) Point partnerships distributions by sex by partnership type						
		# of non-cohabiting partnerships				
		0	1	2	3+	%
Female	# of cohabiting partnerships	0	39.9	28.9	3.4	0.5
		1	26.2	1.0	0.0	0.1
	%	66.1	29.9	3.4	0.6	100.0
Male		0	48.3	20.8	3.9	1.1
		1	20.0	3.3	0.9	0.3
		2	0.9	0.4	0.1	0.0
		%	69.2	24.5	4.9	1.4
					1.4	100.0

(b) Overall point partnership distributions by sex				
	0	1	2	3+
F	39.9	55.1	4.4	0.6
M	48.3	40.8	8.1	2.8

Cell entries in (a) represent proportions of the total sex-specific population (i.e. table totals rather than row totals)

Table 3

Epidemic outcomes

Scenario	Network-transmission pair	Final prevalence in sustained epidemics			# Extinctions before first sustained epidemic		
		Min	Mean ^a	Max	Min	Mean ^b	Max
Baseline	Hollingsworth: baseline network	1.3%	8.8%	12.1%	1	11.2	37
	Hollingsworth: "Rounded up"	12.3%	13.8%	14.6%	1	5.3	10
Stage-specific infectivity variants	High (6.2 coital acts per month during final 5 months)						
	Pinkerton	5.6%	7.2%	9.4%	2	9.0	16
	Wawer	1.4%	3.2%	6.1%	1	9.2	14
	Abu-Raddad	0.2%	1.1%	3.8%	24	73.3	135
	Med. (3.1 coital acts per month during final 5 months)						
	Pinkerton	0.1%	2.0%	6.6%	13	30.8	66
	Wawer	0.3%	1.4%	3.9%	23	60.6	150
	Abu-Raddad	c	c	c	c	c	c
Low (0 coital acts per month during final 5 months)							
	Pinkerton	0.5%	1.0%	1.3%	21	97.6	263
	Wawer	c	c	c	c	c	c
	Abu-Raddad	c	c	c	c	c	c
Counterfactuals	Hollingsworth: no concurrency	c	c	c	c	c	c
	Hollingsworth: no acute peak	c	c	c	c	c	c

^aMean of 5 runs with sustained epidemics in different network realizations

^bMean across 5 different network realizations, with up to 100 runs each

^cNo runs resulted in a sustained for this scenario (1,000 attempted)