Concurrent partnerships and the spread of HIV

Martina Morris and Mirjam Kretzschmar*

Objective: To examine how concurrent partnerships amplify the rate of HIV spread, using methods that can be supported by feasible data collection.

Methods: A fully stochastic simulation is used to represent a population of individuals, the sexual partnerships that they form and dissolve over time, and the spread of an infectious disease. Sequential monogamy is compared with various levels of concurrency, holding all other features of the infection process constant. Effective summary measures of concurrency are developed that can be estimated on the basis of simple local network data.

Results: Concurrent partnerships exponentially increase the number of infected individuals and the growth rate of the epidemic during its initial phase. For example, when one-half of the partnerships in a population are concurrent, the size of the epidemic after 5 years is 10 times as large as under sequential monogamy. The primary cause of this amplification is the growth in the number of people connected in the network at any point in time: the size of the largest 'component'. Concurrency increases the size of this component, and the result is that the infectious agent is no longer trapped in a monogamous partnership after transmission occurs, but can spread immediately beyond this partnership to infect others. The summary measure of concurrency developed here does a good job in predicting the size of the amplification effect, and may therefore be a useful and practical tool for evaluation and intervention at the beginning of an epidemic.

Conclusion: Concurrent partnerships may be as important as multiple partners or cofactor infections in amplifying the spread of HIV. The public health implications are that data must be collected properly to measure the levels of concurrency in a population, and that messages promoting 'one partner at a time' are as important as messages promoting fewer partners.

AIDS 1997, **11**:641–648

Keywords: HIV, networks, sexual behavior, transmission

[For editorial comment, see pp681–683]

Introduction

There is now considerable variation in the timing and intensity of the HIV epidemic in different regions of the world [1]. Explanations for these differences, especially for the rapid and pervasive spread among heterosexual populations in sub-Saharan Africa, have focused on three factors: the rate of sexual partner acquisition [2–5], the impact of certain 'core groups' [6–12], and the presence of other sexually transmitted diseases (STD) that may amplify HIV transmission [13–16]. We investigate the potential impact of another factor: concurrent (or simultaneous) sexual partnerships. Concurrency need not change the rate of partner acquisition, it simply affects the overlap of partners over time. It represents an alternative to sequential monogamy, a different pattern in the general social organization of sexuality rather than a feature of 'core groups'.

Intuition and exploratory work in mathematical disease models [17–20] suggest that concurrent partnerships will amplify the spread of an infectious agent such as

From the Department of Sociology, Pennsylvania State University, University Park, Pennsylvania, USA and the *National Institute of Public Health and Environmental Protection, Bilthoven, The Netherlands.

Sponsorship: Supported by the Rockefeller Foundation, National Institutes of Health (1R29HD34957-01).

Requests for reprints to: Martina Morris, Department of Sociology, Pennsylvania State University, 614 Oswald Tower, University Park PA 16802-6211, USA.

Date of receipt: 6 June 1996; revised: 22 November 1996; accepted: 15 January 1997.

HIV. From the viewpoint of the virus, there is less time lost after transmission occurs in waiting for the current partnership to dissolve, or between the end of one partnership and the beginning of another. In addition, the effect of partner sequence on exposure risk is reduced. Under serial monogamy each subsequent partner increases the risk of infection to a particular subject, so that earlier partners are less likely to be exposed to an infected subject than later partners. If partnerships are concurrent, much of the protective effect of the sequence is lost. Earlier partners remain connected to the subject, and can be exposed when the subject becomes infected by a later concurrent partner.

To formalize and quantify these intuitions we perform a simulation experiment of disease spread under different partnership scenarios. The simulations are based on a fully stochastic model, with a disease transmission process superimposed over a pair-formation and dissolution process in a heterosexual population. The scenarios are carefully structured to ensure that concurrency is not confounded with a simple increase in the number of partnerships. Comparisons across scenarios reflect the difference between, for example, having five partners in sequentially monogamous fashion, and five partners with some or all partnerships overlapping. Although HIV transmission is ultimately the topic of interest, we do not aim to capture the biological complexities of this particular disease (such as asymmetric or variable infectivity). Instead, we employ relatively simple biological assumptions and focus attention on the effects of concurrent partnerships. In the Discussion we will consider how the impact of concurrency might change under alternative assumptions.

Methods

Measuring concurrency

While concurrency is simply defined at the individual level by the number of partnerships a person has at any moment in time, a summary of the amount of concurrency present in a population turns out to be more difficult. We will use a graph theoretic summary index to define and measure concurrency. Although the approach initially appears complicated and abstract, the measure that results ultimately turns out to be quite simple, both to understand, and to estimate from data that are easily collected.

In what follows, we use the terminology of graph theory to describe aspects of the partnership network. The network is called a 'graph', persons are 'nodes', and partnerships are 'links'. The number of simultaneous partners a person has is called 'the momentary degree of a node', and the variation in degree is captured by the degree distribution, mean degree and degree variance in the graph.

The level of concurrency in a population clearly reflects some aspect of the momentary degree distribution. Degree mean alone is not a good measure when nodes can have degree 0, because averaging the links over isolates (persons without partners) and non-isolates obscures the simultaneity of interest in the partnerships. However, restricting the degree mean to non-isolates does not resolve this problem.

An example can be seen in Fig. 1a. The two contact graphs here both have five non-isolate nodes and four links; therefore, they have the same mean degree (whether the mean is based on all nodes, or on non-isolate nodes only). The pattern of concurrency, however, is critically different in the two graphs. One could argue that the second graph (graph II) has 'more' concurrency than graph I, as three nodes have concurrent partnerships in graph II, whereas only one node does in graph I. From an epidemiological perspective, however, graph I has the more efficient concurrency structure: transmission will be more rapid and less variable because the maximum distance between any two nodes in the network is only 2, compared with 4 in graph II.

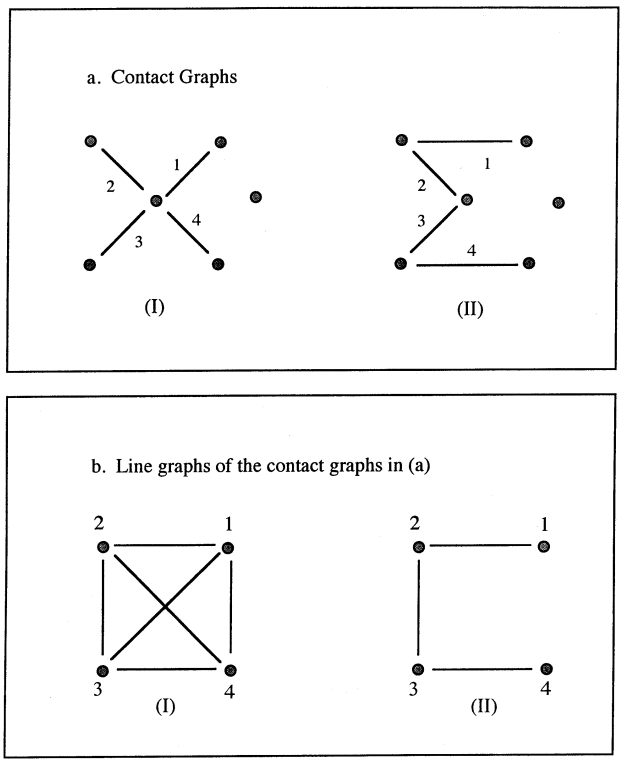

Fig. 1. (a) Two contact graphs with the same number of non-isolate nodes and links, and thus the same mean degree, but different patterns of concurrency. (b) Line graph representations of the two contact graphs. The mean degrees of the line graphs are different, and this measure can be used to represent the epidemiologically relevant aspects of concurrency. See Measuring concurrency (Methods) for further discussion.

To distinguish the epidemiologically relevant aspect of concurrency we define a concurrency measure κ based on the 'line graph' of the original contact graph [21]. A node in the line graph transformation represents a partnership, and an link between two nodes indicates that they share a node in the original graph (i.e., they are concurrent partnerships). An example of the line graph, and how it varies with the different patterns of concurrency is given in Fig. 1b. The two panels are the line graphs of the original contact graphs in Fig. 1a. Every link in Fig. 1a becomes a node in Fig. 1b, and concurrent partnerships in Fig. 1a are connected by a link in Fig. 1b. The degree of a node in the line graph thus represents the number of partnerships that are adjacent to it. [Formally, let G(V,E) be the contact graph (as in Fig. 1a), where the set of nodes V describes the N individuals of the population and the set E (subset of $V \times V$) describes the existing partnerships. The line graph L of G [21] has as nodes the set E. Two nodes of L are linked if they are adjacent edges in G (i.e., they share a node of V).] Under sequential monogamy, the line graph is made up of isolates because no partnership is concurrent with any other. For a star-shaped pattern in the original contact graph, the associated line graph is completely connected. Other patterns fall in between.

We define our population level measure of concurrency κ as the mean degree of the line graph. It represents the average number of concurrent partnerships per partnership in the population. The measure takes the value 0 under sequential monogamy, is equal to 1 when every partnership is concurrent to one other on average, and in general equals the line graph degree mean when the momentary degree distribution is Poisson. Because partnerships are discrete entities, when κ lies between 0 and 1, it has a very simple interpretation and indicates the fraction of partnerships that are concurrent. The measure also has the useful property in that it converges to a simple function of the variance (σ^2) and the mean (μ) of the degree distribution of the original contact graph as the number of nodes grows large [22]:

$$(1) \quad \kappa = \sigma^2/\mu + \mu - 1$$

The importance of this approximation is that it makes the measure easy to estimate from the available data. The approximation depends only on parameters from the original contact graph, so it can be estimated from local, rather than complete, network data [23]. Complete network data, that is, data on all nodes and links in a network, are both expensive and very difficult to collect. When dealing with sexual networks, such data would be virtually impossible to collect routinely. By contrast, local network data, that is, data on randomly selected respondents and their immediate partners only, are relatively easy and inexpensive to collect. One can rely on standard survey sampling techniques, and obtain information on the partner from the respondent so that no contact tracing is necessary. Local network data provide sufficient information for many epidemiological modeling purposes [24]. The fact that such data can also be used to estimate the epidemiologically relevant aspects of concurrency is fortunate, and it dramatically increases the usefulness of the measurement approach presented here.

It is interesting to note that the first two terms of the expression on the right hand side of equation 1 have been proposed as an approximation for the 'effective' mean contact rate when there is significant contact rate variability in the population [25]. This approximation was intended for the accumulation of sequentially monogamous partnerships, to measure variability in the number of partners per person over a period of time, rather than per person at a moment in time. There, as we also find here, variation amplifies the spread of infectious disease.

Simulation model

We simulate the spread of a disease through a population under 10 different scenarios, where the first scenario represents sequential monogamy, and the remaining nine scenarios represent increasing levels of concurrency. We will briefly outline the main assumptions of the model; a complete description of the modeling framework has been published elsewhere [22].

The algorithm for producing the network and spreading the infection involves three separate processes: pair formation, pair dissolution, and disease transmission. Pair-formation and dissolution are controlled by a discrete-time Markov model. New partnerships form with a rate ρ, which depends on the number of partnerships already present, and existing partnerships separate with a constant rate σ, such that the total number of partnerships in the population fluctuates randomly around a constant average. This average is the same for all scenarios, and is held constant to ensure that concurrency is not confounded with a simple increase in the number of partnerships.

The momentary distribution of partnerships (Fig. 1) is independent of ρ and σ, and is varied over 10 levels to represent scenarios ranging from serial monogamy (individuals have no more than one partner at a time), to increasingly higher levels of concurrency (the momentary number of partners per individual is randomly distributed). This is accomplished by conditioning the probability of a new partnership on the current partnership status of each member of the pair. Under sequential monogamy, the probability of a partnership forming between two randomly selected persons is zero if either person has another partner. Under concurrency, this probability is $(1 - \varepsilon)$ where ε ranges from [0–1]

and is decreased to increase the level of concurrency. Mixing between the partnership classes under concurrency is random, that is, individuals with many partners are as likely to pair with monogamous individuals as they are with multipartnered individuals (the impact of assortative and disassortative mixing is analysed in Morris and Kretzschmar [26]). With the parameters chosen here, individuals have a momentary mean of

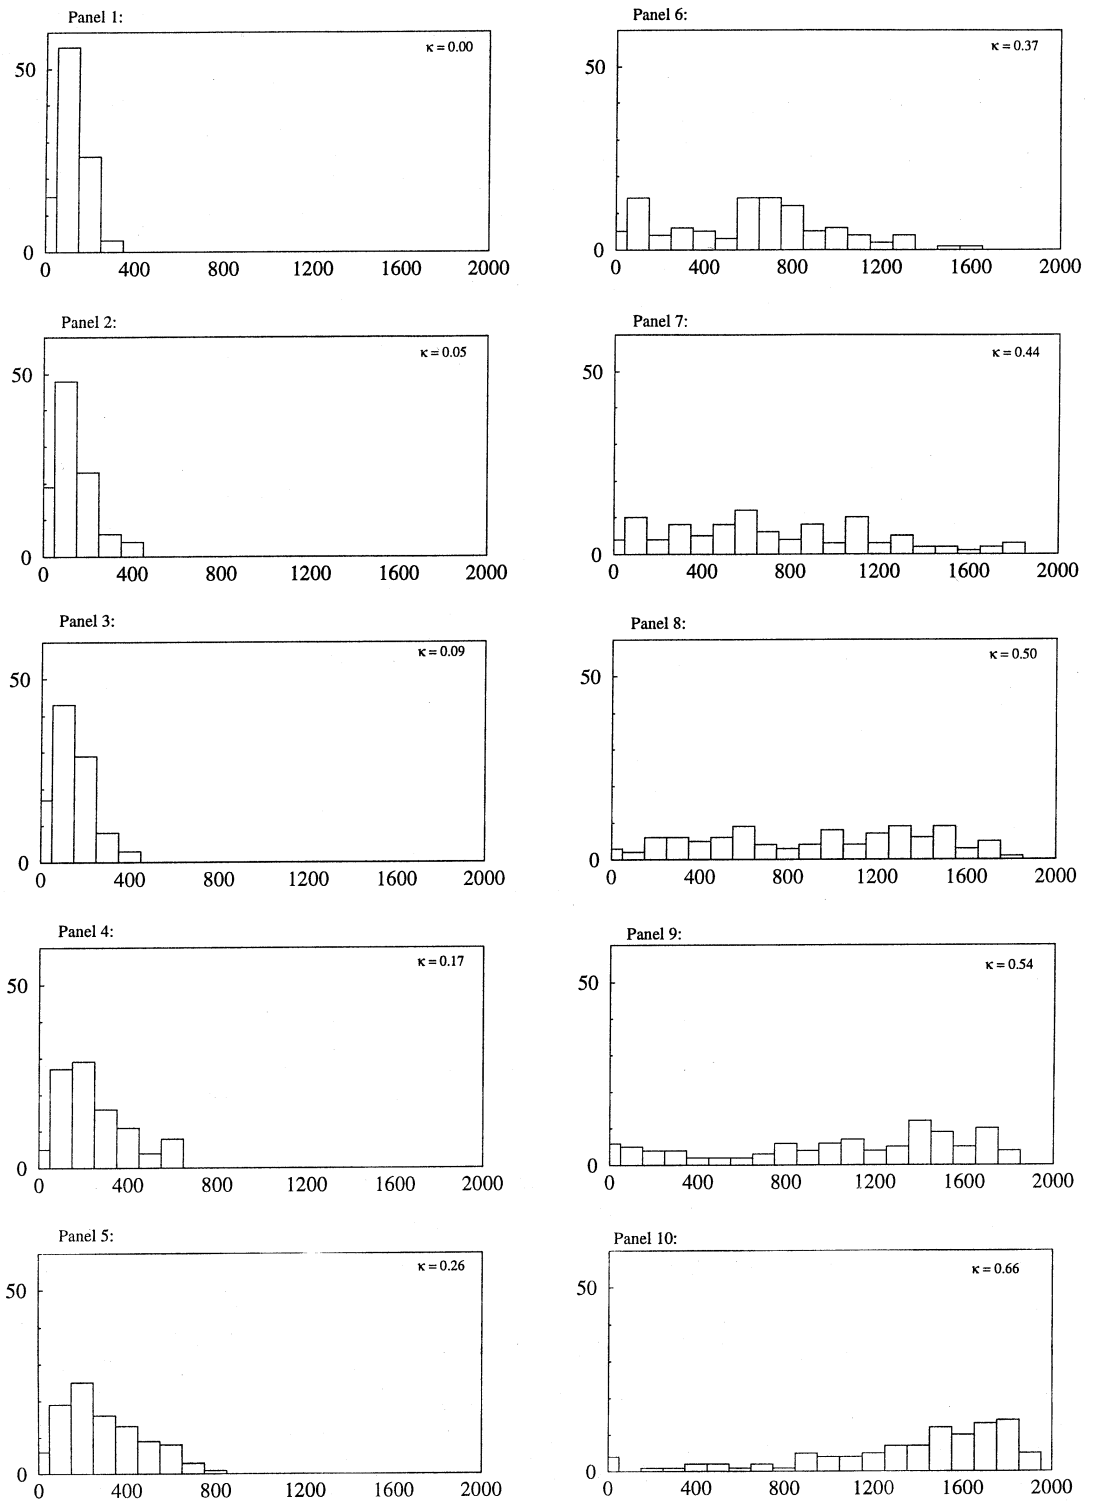

Fig. 2. The distribution of the number of infected cases at the end of the 5-year simulation for each level of concurrency. Each panel represents the results of 100 simulations run at the specified level of concurrency. *x*-axis, number of HIV-infected individuals at the end of the run; *y*-axis, percentage of runs. Panel 1 represents sequential monogamy, and panels 2–10 represent increasing levels of concurrency. The concurrency index, κ, is shown in the upper right corner of each panel. See Measuring concurrency (Methods) for the definition and interpretation of this index.

just under one partner (range, 0–1 for monogamy, 0–4 for concurrency), accumulate an average of about 1.2 partners per year, and the typical partnership lasts about 6–7 months. As measured by κ, the fraction of partnerships that are concurrent ranges from 0 to 0.67. Because the total number of partnerships does not change, increasing the level of concurrency simply reallocates partnerships across the population, causing a change in distribution rather than in magnitude. The distributional change involves both an increase in the number of persons with concurrent partnerships, and a corresponding increase in the number of isolates.

The disease process begins by infecting a randomly selected person in a population with a stationary distribution of partnerships, defined as time $t = 0$. Disease transmission takes place within partnerships between susceptible and infected individuals at a constant rate of 0.05 per day. Each simulation is run for 5 years in a heterosexual population of 2000. For each of the 10 scenarios, 100 simulations are run with the same parameter values in order to observe variation in the epidemic outcomes. There are no vital dynamics or disease-related mortality.

Several aspects of this simulated network and associated epidemic outcomes are measured. The network structure is in dynamic equilibrium, so the summaries of this structure are calculated at the end of each 5-year run. The level of concurrency, κ, is calculated using equation 1 for each run, and the mean of the 100 runs is used for each scenario. The size of the largest component in the network is also recorded at the end of the run, as is the mean used for each scenario. Note that this is a measure of the momentary component size; additional nodes may have been connected over the course of the run but this is not measured here. The size of the epidemic is recorded at the end of the run; both the full distribution and the mean of the 100 runs for each scenario are presented (Figs 2 and 3, respectively). The growth rate is estimated as the change in the log of the mean number of infected persons per unit time during the phase of the epidemic, where the number of persons infected is below $(N/2)^{1/2}$, about 22 in this population (compare Jacquez and O'Neill [27] and Ball and Donnely [28]). Here, again, the mean is taken over the 100 runs for each scenario.

Results

Concurrency dramatically increases both the size and the variability of an epidemic, as can be seen in Fig. 2. Each panel in this figure represents the distribution of the number infected 5 years after the start of the epidemic for the 100 runs under the indicated value of concurrency. The 10 panels represent increasing levels of concurrency, starting with $\kappa = 0$ (serial monogamy) for panel 1, and rising to $\kappa = 0.67$ for panel 10. As concurrency rises, there are two corresponding shifts in the epidemic size distribution. First, the modal size increases. When one-quarter of the partnerships are concurrent ($\kappa = 0.26$, panel 5), the typical epidemic is three times as large by the end of 5 years as it is under sequential monogamy. When one-half of the partnerships are concurrent, the epidemic is about 10 times as large (panel 8). The variability in epidemic size also rises, as can be seen by the increasing spread in the distribution as one moves from panels 1 to 10. Thus, although the epidemic is typically larger, the outcome is also more variable, making it more difficult to predict.

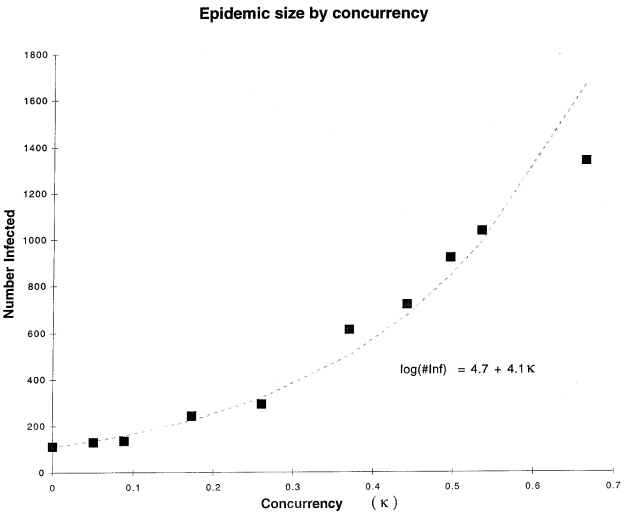

Fig. 3. The mean final size of the epidemic as a function of concurrency. Each observation represents the mean of 100 runs under the same value for the concurrency index κ. The full distribution of epidemic size under each scenario is shown in Fig. 2.

A summary of the relationship between concurrency and epidemic size is shown in Fig. 3. The relationship is approximately linear on the log scale, so the impact of concurrency is quite large: the final number infected increases exponentially with κ. The regression estimate indicates that in this range of κ (i.e., 0–1, where it can be interpreted in percentage terms) for each 10% increase in the average number of concurrent partnerships the final epidemic size rises by a factor of roughly 40%. There are two things to keep in mind about this estimate. The first is that a full unit change in the average number of concurrent partnerships (κ) is a large change: for example, the increase from 0 to 1 represents the change from sequential monogamy to all partnerships being concurrent to one another (on average). Thus, the range examined here may represent most of the range that one would find in real populations. Second, one would expect this relationship to be valid only during the exponential growth phase of the epidemic, as in later phases saturation effects will prevail.

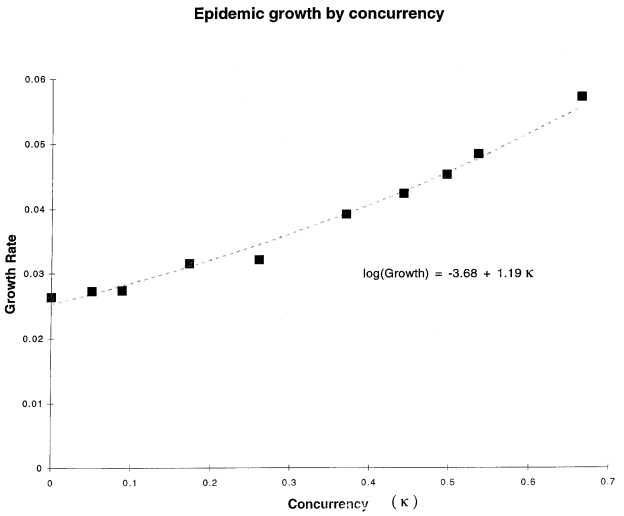

Fig. 4. The mean growth rate of the epidemic as a function of concurrency. Each observation represents the mean of 100 runs under the same value for the concurrency index κ. The growth rate was estimated during the initial 'exponential' phase when the number of infected persons was below $(N/2)^{1/2}$, about 22 in this population. See Simulation model (Methods) for further details.

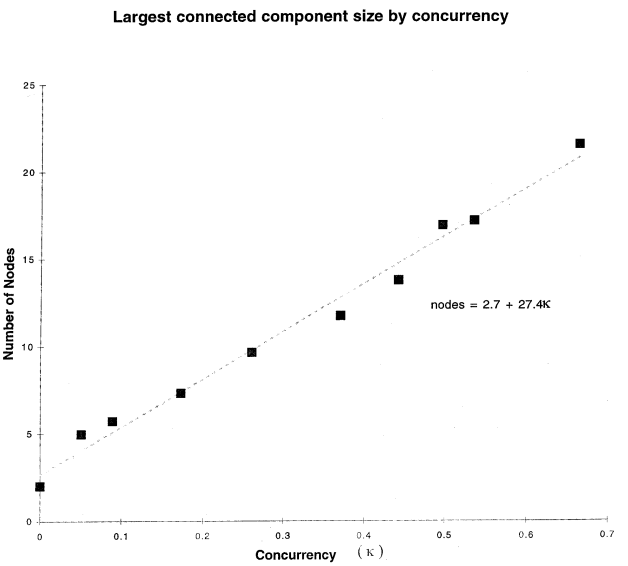

Fig. 5. The typical size of the largest (momentary) connected component as a function of concurrency. Each observation represents the mean of 100 runs under the same value for the concurrency index κ. The size of the largest component size is obtained at the end of each run and averaged over the set of runs for that scenario.

The initial amplification of spread is quite large, however: under the conditions here, it would take over 50 years for the epidemic under monogamy to reach the size observed in 5 years under the higher levels of concurrency (further analysis of changes in the evolutionary path of the epidemic can be found in Morris and Kretzschmar [26]).

Concurrency also increases the speed with which the epidemic spreads. Epidemic growth rate during the initial exponential phase of the epidemic is plotted against κ in Fig. 4. During this phase, the rate of growth also appears to be a exponentially increasing function of κ. The regression estimate indicates that each 10% increase in the average number of concurrent partnerships raises the growth rate by about 12%. This is close to a linear effect.

Although these figures clearly demonstrate that concurrency has a large impact on transmission dynamics, they provide little insight into why. To understand the mechanism operating here, it is necessary to examine how concurrency changes the overall structure of the network. The key aspect of this change appears to be in the effect that concurrency has on increasing the size of the largest 'connected component': the number of nodes in the original graph that are directly or indirectly connected at any point in time. Under sequential monogamy, the maximum size of a connected component can not exceed 2. Under concurrency, by contrast, the maximum size of a connected component can become quite large: individuals have partners who are themselves connected to others, these others are again connected to additional persons, and so on.

The effect of concurrency on the average maximum component size are shown in Fig. 5. Maximum component size ranges here from 2 for $\kappa = 0$ to around 20 for $\kappa = 0.67$ (it rises much higher when partnerships are of longer duration), and the effect of concurrency is roughly linear. This order of magnitude increase in the size of the largest network component explains the dramatic changes in the course of the epidemics simulated: concurrency creates a large, loosely structured, constantly shifting web of connected nodes in a network, enabling an infectious agent to spread rapidly and pervasively.

Discussion

Compared with sequential monogamy, a pattern of concurrent partnerships can dramatically change the early course of an epidemic, increasing the growth rate and the number infected exponentially. These findings are consistent with other recent studies of sexual mixing [24,29] and core-group dynamics [30], which show that the distribution of partnerships, and not simply their mean, is a critical factor for infectious disease transmission. Together, these studies strongly suggest the need for network-oriented survey design and intervention in AIDS prevention efforts.

Although variable infectivity was not used in the models here, one can speculate that the effect of concurrency would be even stronger under that scenario. A short, early peak of infectivity would increase the likelihood that the virus would remain locked in an existing partnership after transmission under sequential monogamy. In contrast, under concurrency the virus can jump across each concurrent connection available during the peak infectious period. If the infectious window is quite short (e.g., a matter of weeks) then fairly high levels of partner change will be needed under sequential monogamy to reach the reproductive threshold and ensure continuing transmission. The larger momentary connected component created by concurrency, however, provides an ideal opportunity for spread, maximizing the use of the transmission window even, and perhaps especially, when partner change rates are low. For the kind of variable infectivity that is hypothesized to characterize HIV, then, the impact of concurrency is likely to be quite high.

The large impact of concurrent partnerships on both epidemic size and speed of spread has two important empirical implications. First, these results provide evidence that concurrent partnerships may partly explain the rapid heterosexual spread of HIV in sub-Saharan Africa and in other populations of the world. Most current research on variation in HIV epidemic intensity is currently focused on the impact of cofactor STD, a focus based on the STD–HIV association documented in observational studies [31,32]. It may be, however, that concurrent partnerships are the root cause of both STD and HIV, and that the STD–HIV correlation is simply an artifact of this process. Better estimates of the attributable risk of STD should become available, as a large, prospective, community-based case–control STD treatment study is now underway in Uganda with a behavioral component designed to measure concurrency [33]. Regardless of the STD effect, however, the simulations here show that concurrent partnerships are an important independent risk factor for HIV transmission. The intervention message is clear: one partner at a time.

The second implication is that patterns of concurrency in populations must be measured if accurate HIV/STD projection and effective intervention are desired. Simple measures of the rate of sexual partner acquisition are clearly inadequate for risk assessment; all of the simulations presented above were based on the same rate of partner acquisition, and the epidemic outcomes varied dramatically. Concurrency and other properties of partnership distribution should be among the core items required for HIV- and STD-related sexual behavior surveys. Such data are not difficult to collect. Many of the important features of the partnership distribution can be obtained using 'egocentric' or local network methods — a standard survey-based approach that relies on representatively sampled respondents to report on the attributes of their partners and relationships [23,34]. This approach is less expensive and intrusive than contact tracing, and thus more feasible for sensitive questions and large populations [24]. Local sexual network studies have now been carried out in several countries around the world, including Uganda [35], Thailand [36], and the United States [37]. The concurrency measure we have presented, κ, can be estimated from such local network data, as it depends only on the mean and variance of the momentary contact distribution.

Finally, there is much to be gained from additional simulation analyses of concurrency. Concurrent partnerships take many forms. The form examined here, of gender symmetric partnerships with relatively short-term partnerships and random mixing, is only loosely based on empirically observed patterns [38,39]. Polygamy features gender asymmetry and disassortative mixing, with women having one partner and men having multiple partners. Another common pattern is a combination of long- and short-term partners, with long-term partnerships formed in a sequential fashion (possibly overlapping during transition), and occasional concurrent short-term partnerships. Non-gender-based heterogeneity could also be modeled, with some persons sequentially monogamous, whereas others have concurrent partnerships. Each of these patterns represents culturally specific aspects of sexual network structure. Some of these structures are more efficient for transmission than others, certain network locations are more at risk than others, and each structure provides different 'weak links' in the chain of transmission that offer opportunities for intervention. HIV-related intervention efforts, whether targeted at risk groups or at overall transmission patterns, would be more effective if such dynamics were known and understood.

References

1. World Health Organization: *The Global Impact of AIDS.* Geneva: WHO; 1992.
2. Serwadda D, Mugerwa R, Sewankambo N, *et al.*: **Slim disease: a new disease in Uganda and its association with HTLV-III infection.** *Lancet* 1985, **ii**:849–852.
3. Piot P, Kreiss JK, Ndinya-Achola JO, *et al.*: **Heterosexual transmission of HIV.** *AIDS* 1987, **1**:199–206.
4. van de Walle E: **The social impact of AIDS in Sub-Saharan Africa.** *Milbank Q* 1990, **68** (suppl 1):10–32.
5. Malamba SS, Wagner H-U, Maude G, *et al.*: **Risk factors for HIV-1 infection in adults in a rural Ugandan community: a case–control study.** *AIDS* 1994, **8**:253–257.
6. D'Costa L, Plummer F, Bowmer I: **Prostitutes are a major reservoir of sexually transmitted diseases in Nairobi, Kenya.** *Sex Transm Dis* 1985, **12**:64–67.
7. Bwayo J, Omari A, Mutere A: **Long distance truck drivers: prevalence of sexually transmitted diseases.** *East Afr Med J* 1991, **68**:425–429.
8. Bwayo J, Plummer F, Omari M, *et al.*: **Human immunodeficiency virus infection in long-distance truck drivers in east Africa.** *Arch Intern Med* 1994, **154**:1391–1396.

9. Kreiss J, Koech D, Plummer F: **AIDS virus infection in Nairobi prostitutes: spread of the epidemic to East Africa.** *N Engl J Med* 1986, **314**:414–418.
10. Miotti PG, Chiphangwi JD, Dallabetta GA: **The situation in Africa.** *Baillieres Clin Obstet Gynaecol* 1992, **6**:165–186.
11. Carswell JW, Lloyd G, Howells J: **Prevalence of HIV-1 in east African lorry drivers.** *AIDS* 1989, **3**:759–761.
12. Orubuloye IO, Caldwell JC, Caldwell P: **The role of high-risk occupations in the spread of AIDS: truck drivers and itinerant market women in Nigeria.** *Int Fam Plann Perspect* 1993, **19**:43–48.
13. Piot P, Tezzo R: **The epidemiology of HIV and other sexually transmitted infections in the developing world.** *Scand J Infect Dis Suppl* 1990, **69**:89–97.
14. Kirby PK, Munyao T, Kreiss J, Holmes KK: **The challenge of limiting the spread of human immunodeficiency virus by controlling other sexually transmitted diseases.** *Arch Dermatol* 1991, **127**:237–242.
15. Wasserheit JN: **Epidemiological synergy. Interrelationships between human immunodeficiency virus infection and other sexually transmitted diseases.** *Sex Transm Dis* 1992, **19**:61–77.
16. Aral SO: **Heterosexual transmission of HIV: the role of other sexually transmitted infections and behavior in its epidemiology prevention and control.** *Annu Rev Public Health* 1993, **14**:451–467.
17. Dietz K, Tudor D: **Triangles in heterosexual HIV transmission.** In *AIDS Epidemiology: Methodological Issues.* Edited by Jewell NP, Kietz K, Farewell VT. Boston: Birkhaeuser; 1992:143–155.
18. Watts CH, May RM: **The influence of concurrent partnerships on the dynamics of HIV/AIDS.** *Math Biosci* 1992, **108**:89–104.
19. Hudson C: **Concurrent partnerships could cause AIDS epidemics.** *Int J STD AIDS* 1993, **4**:349–353.
20. Altmann M: **Susceptible-infected-removed epidemic models with dynamic partnerships.** *J Math Biol* 1995, **33**:661–675.
21. Harary F: *Graph Theory*. Reading: Addison-Wesley; 1969.
22. Kretzschmar M, Morris M: **Measures of concurrency in networks and the spread of infectious disease.** *Math Biosci* 1996, **133**:165–195.
23. Pattison P: *Algebraic Models for Social Networks.* Cambridge: Cambridge University Press; 1993.
24. Morris M: **Data driven network models for the spread of infectious disease.** In *Epidemic Models: Their Structure and Relation to Data*. Edited by Mollison D. Cambridge: Cambridge University Press; 1995:302–322.
25. May RM, Anderson RM: **Transmission dynamics of HIV infection.** *Nature* 1987, **326**:137–142.
26. Morris M, Kretzschmar M: **Concurrent partnerships and transmission dynamics in networks.** *Soc Netw* 1995, **17**:299–318.
27. Jacquez J, O'Neill P: **Reproduction numbers and thresholds in stochastic epidemic models.** *Math Biosci* 1991, **107**:168–186.
28. Ball F, Donnely P: **Branching process approximation of epidemic models.** Presented at the *2nd World Congress of the Bernoulli Society*. Uppsala, 1990.
29. Anderson R, May R, Boily M, Garnett G, Rowley J: **The spread of HIV-1 in Africa: sexual contact patterns and the predicted demographic impact of AIDS.** *Nature* 1991, **352**:581–589.
30. Hethcote H, Yorke JA: *Gonorrhea Transmission Dynamics and Control.* Berlin: Springer Verlag; 1984.
31. Plummer FA, Simonsen JN, Cameron DW, *et al.*: **Cofactors in male–female sexual transmission of human immunodeficiency virus type 1.** *J Infect Dis* 1991, **163**:233–239.
32. Laga M, Manoka A, Kivuvu M, *et al.*: **Non-ulcerative sexually transmitted diseases as risk factors for HIV-1 transmission in women: results from a cohort study.** *AIDS* 1993, **7**:95–102.
33. Wawer MJ: *STD Control for AIDS Prevention, Uganda.* National Institutes of Health research grant RO1-A1 34826, 1993.
34. Burt RS: **Network items and the general social survey.** *Soc Netw* 1984, **6**:293–339.
35. Wawer MJ: *Ugandan Sexual Network/Behaviors Study for HIV Prevention.* National Institutes of Health Research grant, 1993.
36. Wawer MJ: *Behavioral Research for AIDS Prevention in Thailand.* National Institutes of Health Research grant, 1990.
37. Laumann E, Gagnon J, Michael R, Michaels S: *The Social Organization of Sexuality*. Chicago: University of Chicago Press; 1994.
38. Larson A: **Social context of human immunodeficiency virus transmission in Africa: historical and cultural bases of East and Central African sexual relations.** *Rev Infect Dis* 1989, **11**:716–731.
39. Obbo C: **HIV Transmission through social and geographical networks in Uganda.** *Soc Sci Med* 1993, **36**:949–955.