Concurrency can drive an HIV epidemic by moving R_0 across the epidemic threshold

Ka Yin Leung[a,b] and Mirjam Kretzschmar[b,c]

Objective: The objective of this study is to investigate whether concurrency can drive an HIV epidemic by moving R_0 across the epidemic threshold.

Design and methods: We use a mathematical framework for a dynamic partnership network and the spread of a one-stage infection to study how concurrency is related to the basic reproduction number R_0. Two concurrency indices were used to measure the level of concurrency. The model allows varying the level of concurrency in the population, while other key network properties such as partnership duration and lifetime number of partners are kept fixed. In this way, the effect of concurrency on R_0 is investigated as an isolated phenomenon.

Results: We find that an increase in concurrency is associated with an increase of R_0. For plausible parameter sets for MSM populations, R_0 is always above the epidemic threshold of 1. For scenarios that are plausible for sub-Saharan African populations, we show that increasing the level of concurrency can lead to R_0 crossing the epidemic threshold. This occurs already at low levels of concurrency. Only a slight shift of the network structure from a purely monogamous population to one wherein individuals are allowed to have at most two partners is enough for this to happen.

Conclusion: Concurrency can be a driver of an HIV epidemic in sub-Saharan Africa for low levels of concurrency, although it is not decisive in MSM populations. A small increase in the level of concurrency can lead to R_0 crossing the epidemic threshold in a sub-Saharan African setting.

AIDS 2015, **29**:1097–1103

Keywords: concurrent partnerships, HIV, mathematical model, R_0, sub-Saharan Africa

Introduction

In the past, there has been much discussion on the role of concurrent partnerships in driving the HIV epidemic in heterosexual sub-Saharan African (SSA) populations. Establishing conclusive evidence of the role of concurrency in empirical investigations is difficult [1–3]. One of the reasons is the fact that concurrency cannot be studied as an isolated phenomenon, but is connected to sexual behaviour in all its complexity and is interwoven with many other network structural properties [4]. More specifically, when the distribution of partnerships over a population or the duration of partnerships change, it will also change the fraction of concurrent partnerships that are observed at any time or accumulated over a longer time period.

On the contrary, mathematical models consistently show the impact of concurrency on incidence and prevalence of HIV [5–9]. Watts and May [5] were one of the first to use a mathematical model to investigate the influence of concurrency on HIV dynamics. The authors used an intuitive notion of concurrency and focused on the dynamics over time for R_0 more than 1. They found that concurrency could lead to a very fast initial spread compared with a monogamous population. Others have mainly used simulation models to investigate the relation between concurrency and initial growth rate or endemic

[a]Mathematical Institute, Utrecht University, [b]Julius Center for Primary Care and Health Sciences, University Medical Center Utrecht, and [c]Center for Infectious Disease Control, National Institute for Public Health and the Environment, Utrecht, The Netherlands.

Correspondence to Ka Yin Leung, Mathematical Institute, Utrecht University, Budapestlaan 6, 3584 CD Utrecht, The Netherlands.
E-mail: k.y.leung@uu.nl
Received: 6 January 2015; revised: 12 March 2015; accepted: 13 March 2015.

DOI:10.1097/QAD.0000000000000676

ISSN 0269-9370

situations [6–9]. In such models, it is inherently difficult to control for different sexual behaviour properties. Also, earlier modelling results have been refuted, because they used unrealistic parameter values and neglected demographic flow through the population [10].

In ref. [11], it was suggested that for a factor to be considered a driver of HIV, one needs to show that this factor can drive the basic reproduction number R_0 across the threshold value of 1, although all other factors remain unchanged. Using a flexible mathematical model in which different sexual behaviour properties can easily be kept fixed, we investigated whether concurrency can be a driver of HIV in this way.

Here, we focused on a model of a one-stage infection (susceptible – infected) without recovery [12,13]. In this model, formation and separation of partnerships is dynamic, and demographic flow through the sexually active population is incorporated. It allows a change in the concurrency level in a population, although other key parameters describing the contact network are kept fixed. One of the central notions of the model that distinguishes it from earlier approaches is the so-called 'partnership capacity', a number n that denotes the maximum number of partners that an individual can have at any time. Pair-formation models for monogamous populations implicitly take $n=1$. Infection can then be transmitted in a partnership between an infectious and susceptible individual. The model does not assume any disease-related mortality. For this model, it is possible to derive an explicit expression for the basic reproduction number R_0. This has the advantage that R_0 can easily be studied as functions of infection and sexual behaviour and parameters.

We investigated how R_0 depends on concurrency when duration of partnerships and lifetime number of partners are kept constant. The parameter choices are based on estimates from existing literature. We used a measure for concurrency that allows us to compare populations with different levels of concurrency. The concurrency measure was then varied by varying partnership capacity n only. Our aim was not to provide a realistic estimate of how HIV transmission depends on concurrency in specific populations, but to provide a proof of principle that concurrency alone can be a driving factor of HIV transmission by moving the reproduction number across the threshold value of 1.

Materials and methods

The method used in this article is as follows; we have a model for the spread of infection in a concurrent partnership network wherein mean sexually active lifetime $1/\mu$, mean partnership duration d_P, mean lifetime number of partners θ and transmission rate β are kept fixed. Concurrency is purely a network characteristic and can be measured using concurrency indices. The basic reproduction number R_0 can be calculated for this model. The relation between concurrency and R_0 is then investigated for six scenarios. The different aspects of our method are explained below (a more detailed description is given in the Supplemental Digital Content).

Partnership network

In ref. [12], a model was introduced for a dynamic partnership network with demographic flow. In this model, we consider a population in which we do not distinguish between sexes and all individuals have the same characteristics, that is the population is homogeneous. The model generalizes earlier pair-formation models to a situation wherein concurrent partnerships are possible. This is done by introducing a maximum number of partners an individual can have simultaneously. This number n is called the partnership capacity. Monogamous individuals then have partnership capacity $n=1$. One may think of an individual as having n 'binding sites' for partners that are either free or occupied by a partner. The crucial (but also limiting) assumption is that these binding sites behave independently of one another as far as partnership formation and separation is concerned.

Concurrency

Concurrency is related to many other sexual behaviour characteristics. The most straightforward of these are partnership duration and lifetime number of partners. Indeed, the longer the duration of a partnership, the more overlap there can be with other partnerships and the more partners an individual has, the more it can have at the same time. Therefore, in order to be able to study the effect of concurrency on infection dynamics, these quantities need to be kept fixed when varying the level of concurrency.

Our simple model has only four parameters. We can fix mean sexually active lifetime $1/\mu$, mean partnership duration d_P and mean lifetime number of partners θ. Then, the level of concurrency is varied by varying partnership capacity n. This parameter is a theoretical construct that does not have any measurable variable associated with it. Therefore, for fixed network properties $1/\mu$, d_P and θ, we quantified concurrency in terms of concurrency indices.

Various ways of quantifying concurrency in a population have been suggested, both in theoretical context [12,14] and for practical use [15,16]. Here, we use the partnership-based concurrency index κ_P introduced in [12] and compare it to the point prevalence of concurrency as suggested by UNAIDS [15] that we denote here by κ_U. The point prevalence of concurrency is the fraction of the population with two or more partners, while the partnership-based concurrency index κ_P can be interpreted as follows: choose a partnership at random and consider

one of the two individuals in this partnership. κ_P then describes the number of other partners of that individual. In the model, κ_P can be computed explicitly [12].

Infectious disease dynamics

In the partnership network, we considered the spread of an infection without recovery as described in ref. [13]. We assumed that infection does not impact partnership formation or separation nor does it impact the demographic process. An infectious individual transmits infection to its susceptible partner at a rate β. A susceptible individual becomes infectious at the very instant that it gets infected and stays infectious with the same level of infectiousness for the rest of its sexual lifetime. For this model, we characterized the basic reproduction number R_0 [13] that we can evaluate numerically for given values of model parameters $1/\mu$, d_P, θ, n and β.

Definition of scenarios

To study the impact of concurrency in different types of populations, we defined a number of scenarios. These descriptions are caricatures only. The model considers a homogenous population. Therefore, it is fully determined by mean sexually active lifetime $1/\mu$, mean partnership duration d_P, mean lifetime number of partners θ, partnership capacity n and transmission rate β.

We kept $1/\mu = 40$ years [17] and $\beta = 0.12$ per year [18,19] fixed throughout this investigation. A scenario is defined by one set of parameter values for the partnership duration d_P and lifetime number of partners θ (Table 1). For each scenario, we considered populations with different concurrency levels by varying partnership capacity n. The scenarios are based on estimates from literature as follows.

We defined four scenarios representing sexual behaviour in SSA. Because of the large discrepancies between male and female sexual behaviour as reported in the literature, we considered separate scenarios on the basis of male and female respondents in SSA heterosexual populations. Furthermore, we also defined distinct scenarios on the basis of spousal versus nonspousal partnerships. Here, SSA 1 and SSA 2 represent heterosexual SSA women in their spousal and nonspousal behaviour, while SSA 3 and SSA 4 represent the spousal and nonspousal behaviour of heterosexual SSA men (for order of magnitude of these values, we used [16,20–23]). Finally, we defined two scenarios meaning to represent sexual behaviour of MSM populations in their steady and casual partnerships, denoted by MSM 1 and MSM 2, respectively [24,25]. The six scenarios are presented in Table 1. The parameter value estimates are explained in more detail in the Supplemental Digital Content.

Table 1. The six scenarios.[a]

Scenario	Partnership duration d_P (years)	Lifetime number of partners θ
SSA 1 (women spousal)	10	2
SSA 2 (women nonspousal)	2.4	4.5
SSA 3 (men spousal	5	3
SSA 4 (men nonspousal)	1.7	6
MSM 1	1.5	45
MSM 2	0.083	135

[a]Six scenarios defined by mean partnership duration d_P and mean lifetime number of partners θ. Scenarios SSA 1–4 represent spousal and nonspousal behaviour of men and women in sub-Saharan Africa and MSM 1,2 represent MSM populations. The parameter choices are based on estimates from published literature such that the orders of magnitudes are within plausible range.

Although the parameter choices are based on estimates from published literature, these scenarios are not meant to represent the complexity of sexual behaviour in populations nor did we aim to make precise/detailed quantitative statements for those populations. We chose parameters such that their orders of magnitude are within a plausible range. Our aim was to show that R_0 can cross the threshold value of 1 when only the level of concurrency is increased.

In the section below, we first studied the structure of the sexual network. In particular, we computed degree distributions and concurrency indices. We then studied the dependency of R_0 on model parameters and concurrency indices.

Results

Sexual network and concurrency

First, we considered the network by considering degree distributions (where the degree of an individual is the number of partners it has) for different values of partnership duration d_P and lifetime number of partners θ and for partnership capacities $n = 3$ and $n = 6$; see Fig. 1.

The degree distribution may look very similar for different partnership capacities (Fig. 1a, c). Note, however, that they are never equal. There is always a nonzero fraction of the population with the maximum number of n partners. For instance, in a population with $n = 6$, there is always a small fraction of the population with six partners, while for $n = 3$, no individual has more than three partners (by definition of n). That small fraction with many concurrent partners plays an important role for transmission. Therefore, a shift of n to larger values can have a significant impact even if the degree distribution seems hardly affected. We see that when n increases, the degree mean does not change (degree mean is equal to the product $\theta d_P \mu$, which is independent of n), while the variance does (Table 2). However, the variance is almost always smaller than the mean, that is the degree distribution is underdispersed. Compared with observed degree distributions (see e.g. [9]), the means in Table 2 are either lower or higher. We

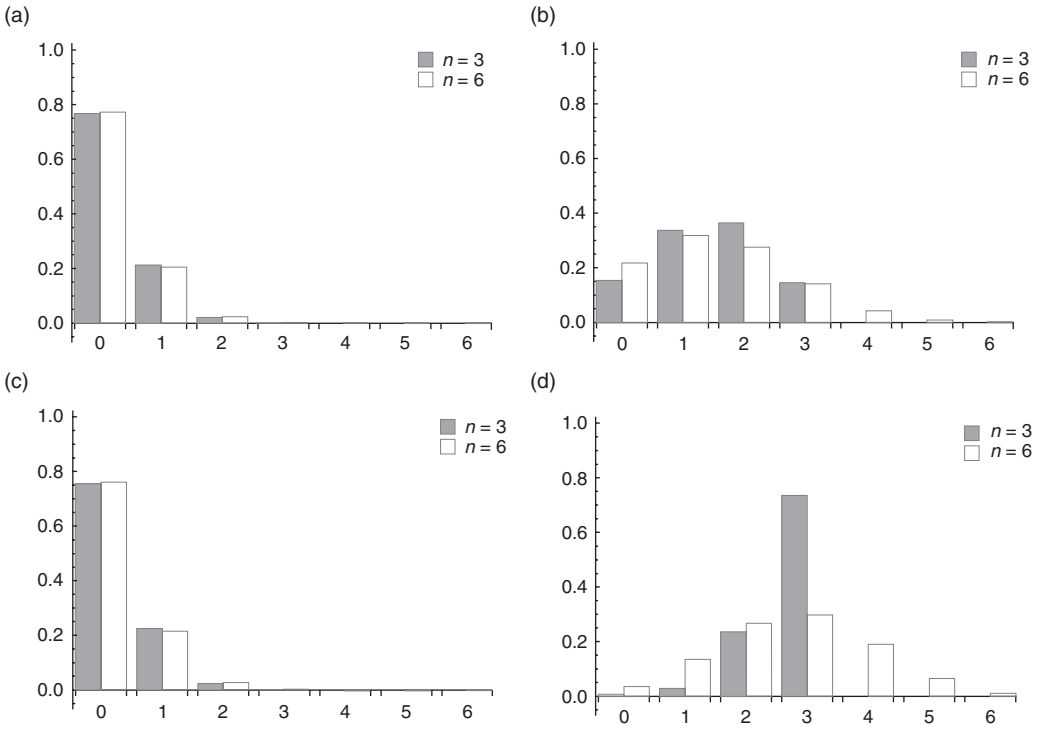

Fig. 1. Degree distributions for different parameter values. d_P denotes mean partnership duration and θ mean lifetime number of partners. The parameter values for (a) and (c) correspond to scenarios SSA 4 and SSA 2, respectively (see also Table 1). (a) $d_P = 1.7$ years, $\theta = 6$ (SSA 4). (b) $d_P = 10$ years, $\theta = 6$. (c) $d_P = 2.4$ years, $\theta = 4.5$ (SSA 2). (d) $d_P = 2.4$ years, $\theta = 45$.

expect the variance to be larger than the mean in real populations.

We next studied how the concurrency indices κ_P and κ_U depend on the parameters of the sexual network (Fig. 2). Intuitively, concurrency measures should increase as a function of partnership duration d_P (when keeping n and θ fixed) and as a function of θ (when keeping n and d_P fixed). Indeed, the longer partnership durations are or the higher the lifetime numbers of partners, the more overlap there has to be in order to 'fit' all partnerships into an individual's life. We found that this indeed holds for κ_P and κ_U. Using its explicit expression (see [12] and Supplementary Digital Contents for details), we see that κ_P is strictly increasing in d_P and θ (Fig. 2a,b). Numerical investigation of κ_U showed that also this is strictly increasing in d_P and θ (Fig. 2c,d).

Table 2. Mean and variance for different parameter values.[a]

Parameter values	n	Mean	Variance
$d_P = 1.7$ years, $\theta = 6$ (SSA4)	3	0.26	0.23
	6	0.26	0.25
$d_P = 10$ years, $\theta = 6$	3	1.5	0.85
	6	1.5	1.32
$d_P = 2.4$ years, $\theta = 4.5$ (SSA2)	3	0.27	0.25
	6	0.27	0.29
$d_P = 2.4$ years, $\theta = 45$	3	2.7	0.28
	6	2.7	1.59

[a]Sexual network characteristics captured by the mean and variance for different parameter value combinations of mean partnership duration d_P, mean lifetime number of partners θ and partnership capacity n. The first and third set of parameter values correspond to scenarios SSA 4 and SSA 2, respectively (indicated between brackets; see also Table 1 for the scenarios).

We also expect that, for fixed partnership duration d_P and lifetime number of partners θ, an increase in n yields an increase in the concurrency indices. Indeed, the larger the maximum number n of simultaneous partnerships individuals may have, the more concurrent partnerships there may be in the population. We immediately find that this is the case for κ_P (using its explicit expression; see [12] and Supplementary Digital Contents for details). However, this is not necessarily the case for κ_U as we find in Fig. 2 c,d. Here, we see that κ_U may decrease when n increases. The reason for this, perhaps counterintuitive, result becomes clear when looking at the changes in degree distribution with increasing n for those situations (see Fig. 1d). With increasing partnership capacity, some individuals will have more partners, but at the expense of a larger proportion of single individuals in the population. The index κ_U is sensitive only to the latter, not to the former. So, it could be that the fraction of the population that has more than one partner decreases, whereas the total number of overlapping partnerships increases. We concluded that κ_U is not useful for our analysis and from here on focus on κ_P.

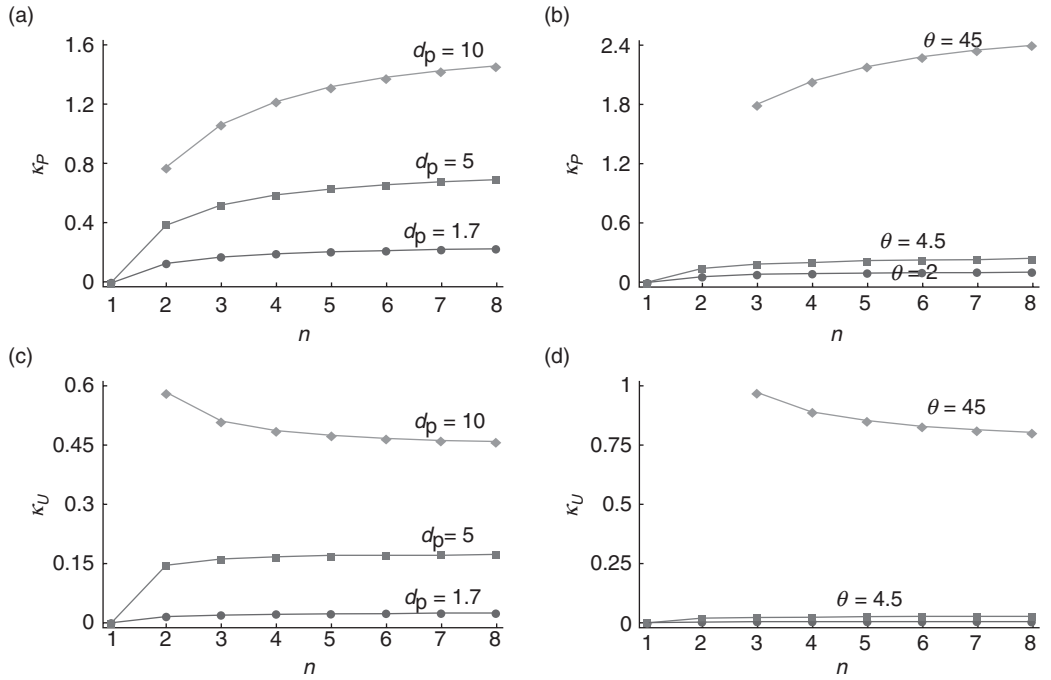

Fig. 2. Concurrency indices κ_P and κ_U as functions of partnership capacity n for different values of mean partnership duration d_P and mean lifetime number of partners θ. (a) $\theta = 6$; (b) $d_P = 2.4$ years; (c) $\theta = 6$; (d) $d_P = 2.4$ years.

R_0 as a function of partnership capacity n

In the same way as we did for κ_P and κ_U in the previous section, we also studied R_0. Using the explicit expression for R_0 (see [13] and Supplementary Digital Content for details), we found that R_0 is increasing in n. Therefore, we can immediately conclude that the concurrency index κ_P is always positively correlated to R_0 when varying n only. In what follows, we studied this in more detail, in particular, we were interested in R_0 for the six scenarios.

Impact of concurrency on R_0

We considered how κ_P and R_0 are related for the six scenarios (Table 1). The results are presented in Fig. 3.

We see that an increase in n alone yields an increase in both κ_P and R_0. We also find that the MSM scenarios yield relatively large R_0-values compared with the SSA scenarios (Fig. 3a). In particular, for the smallest n, R_0 is already larger than 1. It also does not seem to be very sensitive to changes in the concurrency level. Note that MSM 1 yields larger values for R_0 than MSM 2. Even though the lifetime number of partners in MSM 1 is much smaller, the mean partnership duration in MSM 2 is much shorter. Therefore, the probability of transmission within a partnership is also much smaller in MSM 2.

Next, we considered the SSA scenarios in Fig. 3b. For all four SSA scenarios, we found that R_0 crosses the threshold

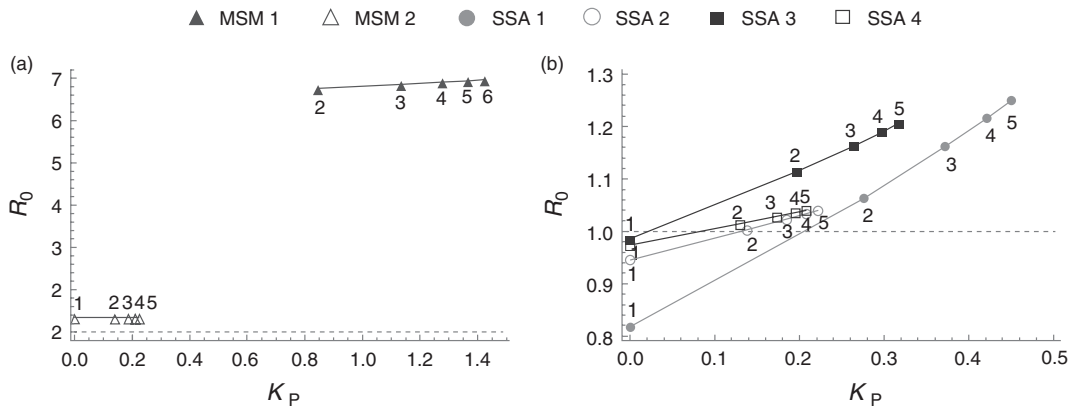

Fig. 3. The impact of concurrency on R_0 by varying partnership capacity n for the six scenarios. The six scenarios are presented in Table 1. Numbers at each point reflect the corresponding partnership capacity number n. (a) MSM scenarios (b) SSA scenarios.

value of 1 when concurrency is increased by increasing n. This shows that it is possible that a population wherein at first no epidemic outbreak is possible can change into a population wherein an epidemic outbreak is possible when more partnerships are concurrent (although all other sexual behaviour remains unchanged). So, we see that concurrency can be a driver of HIV for these SSA scenarios. The crossing of the epidemic threshold occurs already at low levels of concurrency (between $n = 1$ and $n = 2$ so for small values of n).

The sensitivity analysis in the Supplemental Digital Content shows that whether or not R_0 crosses the threshold value of 1 depends on the parameter values. It is certainly not always the case.

Discussion and conclusion

Using a model for a dynamic sexual network and the transmission of infection, we analysed the association between concurrency and R_0. We found that, for parameter sets that are plausible for MSM populations, R_0 is always above 1. Therefore, concurrency is not a driver of HIV in these MSM populations in the sense that in these populations, other factors such as the lifetime number of partners are more important in determining R_0. We also see that concurrency does not really have an effect on the two MSM scenarios, supporting our conclusion that it is not a driver in these MSM populations. However, for parameter sets that are plausible for SSA populations, we found that an increase of the level of concurrency in a population can drive the basic reproduction number from below to above the threshold of 1. Only a slight shift of the network structure from a purely monogamous population to one wherein individuals have at most two partners is required for this to happen.

Furthermore, we showed that two widely used concurrency indices have different properties. In particular, although the partnership-based concurrency index κ_P is always increasing with increasing partnership capacity n, this is not necessarily the case for the point-prevalence of concurrency κ_U. The latter can possibly decrease with increasing n, leading to some doubt about the usefulness of κ_U in measuring the impact of concurrency on transmission of infection.

The level of concurrency as measured by the index κ_P is closely related to other network properties such as the lifetime number of partners and partnership duration. By defining a model in which all parameters except for the partnership capacity n can be kept at a fixed value, we were able to study the impact of changing concurrency levels without changing lifetime numbers of partners and partnership duration. As the network is dynamic, the number of partners at age a are changing due to partnership formation and separation over the course of an individual's life, even if lifetime number of partners is defined as a constant parameter.

In all scenarios, we used the same estimate for the transmission rate β. However, it is expected that the transmission rate is higher in MSM populations than in heterosexual populations due to higher risk in sexual practices and higher rates of sexual acts within partnerships; see for example [26] for estimated per-act probabilities for different exposure routes. A larger transmission rate β for MSM scenarios would also yield a larger R_0 and support our conclusion that concurrency is not decisive for the epidemic in those populations.

Our results also illustrate that that the value of R_0 is not merely determined by κ_P, but other parameters (partnership duration d_P, lifetime number of partners θ and transmission rate β) play an important role as well (compare e.g. scenarios SSA 2 and SSA 4 to MSM 2 in Fig. 3, which all have similar κ_P values but different R_0 values).

Finally, we have only investigated whether concurrency can lead to R_0 crossing the threshold value 1 and therefore allowing for an epidemic outbreak. In other words, we have only investigated the impact of concurrency on the beginning of an epidemic, not its impact on a mature epidemic that has reached its endemic state. The role of concurrency for persistence of HIV in an endemic state could be different and has to be investigated in a different way. Preliminary numerical investigations with our model have shown that concurrency is also positively correlated with endemic prevalence for the sub-Saharan African scenarios. However, this has to be investigated further and is outside the scope of this article.

Our modelling approach has a number of limitations that have to be kept in mind. First, we considered a partnership network in a homogeneous population with only one type of individuals and one type of partnerships. The model at present neglects all population heterogeneity and differences between men and women. However, our framework can be generalized to a two-sex population with asymmetry between men and women in partnership capacity. Furthermore, our framework can be modified to allow for two (or multiple) types of partnerships. By distinguishing different types of partnerships, for example spousal and nonspousal, it is possible to take different partnership dissolution rates into account. See ref. [12] for a more detailed description on various generalizations of the model.

Next, we do not take any disease-related mortality into account. In the case of disease-related mortality, an infectious individual has a shorter life expectancy and less secondary cases may occur than in the case without disease mortality, that is the R_0 will be smaller. This means that, with the addition of disease mortality, R_0 will then

cross the threshold value at somewhat higher levels of concurrency than shown in Fig. 3. The addition of disease-related mortality will not change the mechanism demonstrated in this article.

The most important limiting assumption, however, is that partnerships are considered to be independent of each other. In other words, whether an individual already has a partner or not has no influence on his propensity to acquire another partner. Second, infection does not influence partnership dynamics nor does it lead to an increased mortality. Finally, we assumed that infectivity is constant throughout the infectious period, thereby neglecting the possibility of high transmission rates during early HIV infection.

Although these are serious limitations when one wants to make quantitative prediction in a real population, our aim here was to establish the possible role of concurrency as a factor in driving the basic reproduction number R_0 across the threshold value of 1 without changing other network properties. We conclude that not only is this possible in a dynamic sexual network but also that for parameter sets that are plausible for heterosexual SSA populations, crossing the threshold occurs for a shift of concurrency from monogamous to very low levels of concurrency.

Acknowledgements

We thank O. Diekmann for helpful input and discussions on the study design and L. Hermans and two anonymous reviewers for helpful comments on the article. Part of the research reported here was carried out while K.Y.L. participated in IIASA's 2013 Young Scientists Summer Program. She conveys thanks to J.A.J. Metz, U. Dieckmann and R. Mazzucco for helpful advice. K.Y.L. is supported by the Netherlands Organisation for Scientific Research (NWO) through grant Mozaïek 017.009.082.

Both authors contributed to the study design, discussion of the results and manuscript draft. K.Y.L. conducted the numerical investigation. Both authors have read and approved the text as submitted to *AIDS*.

Conflicts of interest

There are no conflicts of interest.

References

1. Lurie MN, Rosenthal S. **Concurrent partnerships as a driver of the HIV epidemic in sub-Saharan Africa? The evidence is limited.** *AIDS Behav* 2009; **14**:17–24.
2. Mah TL, Halperin DT. **Concurrent sexual partnerships and the HIV epidemics in Africa: evidence to move forward.** *AIDS Behav* 2010; **14**:11–16.
3. Tanser F, Bärnighausen T, Hund L, Garnett GP, McGrath N, Newell ML. **Effect of concurrent sexual partnerships on rate of new HIV infections in a high-prevalence, rural South African population: a cohort study.** *Lancet* 2011; **378**:247–255.
4. Kretzschmar M, White RG, Caraël M. **Concurrency is more complex than it seems.** *AIDS* 2010; **24**:313–315.
5. Watts CH, May RM. **The influence of concurrent partnerships on the dynamics of HIV/AIDS.** *Math Biosci* 1992; **108**:89–104.
6. Morris M, Kretzschmar M. **Concurrent partnerships and the spread of HIV.** *AIDS* 1997; **11**:641–648.
7. Eaton JW, Hallett TB, Garnett GP. **Concurrent sexual partnerships and primary HIV infection: a critical interaction.** *AIDS Behav* 2011; **15**:687–692.
8. Kim J-H, Riolo RL, Koopman JS. **HIV transmission by stage of infection and pattern of sexual partnerships.** *Epidemiology* 2010; **21**:676–684.
9. Goodreau SM, Cassels S, Kasprzyk D, Monta DE, Greek A, Morris M. **Concurrent partnerships, acute infection and HIV epidemic dynamics among young adults in Zimbabwe.** *AIDS Behav* 2012; **16**:312–322.
10. 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.
11. Kretzschmar M, Caraël M. **Is concurrency driving HIV transmission in sub-Saharan African sexual networks? The significance of sexual partnership typology.** *AIDS Behav* 2012; **16**:1746–1752.
12. Leung KY, Kretzschmar MEE, Diekmann O. **Dynamic concurrent partnership networks incorporating demography.** *Theor Popul Biol* 2012; **82**:229–239.
13. Leung KY, Kretzschmar MEE, Diekmann O. **SI infection of a dynamic partnership network: characterization of R0.** *J Math Biol* 2014; doi:10.1007/s00285-014-0808-5.
14. Kretzschmar M, Morris M. **Measures of concurrency in networks and the spread of infectious disease.** *Math Biosci* 1996; **133**:165–195.
15. UNAIDS Reference Group on Estimates, Modelling, and Projections: Working Group on Measuring Concurrent Sexual Partnerships. **HIV: consensus indicators are needed for concurrency.** *Lancet* 2010; **375**:621–622.
16. Glynn JR, Dube A, Kayuni N, Floyd S, Molesworth A, Parrott F, *et al*. **Measuring concurrency: an empirical study of different methods in a large population-based survey in northern Malawi and evaluation of the UNAIDS guidelines.** *AIDS* 2012; **26**:977–985.
17. World Bank. Life expectancy at birth. http://data.worldbank.org/indicator/SP.DYN.LE00.IN/countries/ZF?display=graph, 2011. [Accessed 27 March 2014].
18. Hollingsworth TD, Anderson RM, Fraser C. **HIV-1 transmission, by stage of infection.** *J Infect Dis* 2008; **198**:687–693.
19. Wawer M, Gray R, Sewankambo N, Serwadda D, Li X, Laeyendecker O, *et al*. **Rates of HIV-1 transmission per coital act, by stage of HIV-1 infection, in Rakai, Uganda.** *J Inf Dis* 2005; **191**:1403–1409.
20. Porter L, Hao L, Bishai D, Serwadda D, Wawer M, Lutalo T, *et al*. **HIV status and union dissolution in sub-Saharan Africa: the case of Rakai, Uganda.** *Demography* 2004; **41**:465–482.
21. Helleringer S, Kohler HP. **Sexual network structure and the spread of HIV in Africa: evidence from Likoma Island, Malawi.** *AIDS* 2007; **21**:2323–2332.
22. Morris M, Epstein H, Wawer M. **Timing is everything: international variations in historical sexual partnership concurrency and HIV prevalence.** *PLoS One* 2010; **11**:e14092.
23. Powers KA, Hoffmann IF, Ghani AC, Hosseinipour MC, Pilcher CD, Price MA, *et al*. **Sexual partnership patterns in Malawi: implications for HIV/STI transmission.** *Sex Transm Dis* 2011; **38**:657–666.
24. Xiridou M, Geskus R, de Wit J, Coutinho R, Kretzschmar M. **The contribution of steady and casual partnerships to the incidence of HIV infection among homosexual men in Amsterdam.** *AIDS* 2003; **17**:1029–1038.
25. Glick SN, Morris M, Foxman B, Aral SO, Manhart LE, Holmes KK. **A comparison of sexual behavior patterns among men who have sex with men and heterosexual men and women.** *J Acquir Immune Defic Syndr* 2012; **60**:83–90.
26. Patel P, Borkowf CB, Brooks JT, Lasry A, Lansky A, Mermin J. **Estimating per-act HIV transmission risk.** *AIDS* 2004; **28**:1509–1519.