Journal of Economic Interaction and Coordination (2019) 14:293–315
https://doi.org/10.1007/s11403-018-00234-1

REGULAR ARTICLE

Partnership duration and concurrent partnering: implications for models of HIV prevalence

Alan G. Isaac[1] · **Larry Sawers**[1]

Received: 27 December 2017 / Accepted: 31 December 2018 / Published online: 10 January 2019

Abstract
Researchers and policy makers have argued that long-duration concurrent relationships promote the spread of HIV. The concurrency hypothesis proposes that concurrent partnering, particularly as manifested in formal and informal polygyny, is a primary contributor to the spread of HIV in sub-Saharan Africa. We investigate claims that agent-based models of concurrent partnering support this hypothesis. Specifically, we explore how assumptions about the duration and network structure of sexual partnerships affect the results of agent-based models of HIV propagation. We offer new support for the contention that long-duration concurrent partnering can be protective against HIV transmission rather than promoting it. Additionally, we argue that the focus on concurrency has misdirected attention away from the key role of exclusivity.

Keywords Concurrency · HIV · Sub-Saharan Africa · Partnership duration · Coital dilution · Exclusivity · Pair formation

JEL Classification I12 · I18 · C63

1 Introduction and background

Although the AIDS epidemic peaked in the late 1990s, tens of millions of people remain infected. There are almost two million AIDS-related deaths annually, which are slightly outpaced by new HIV infections. Eastern and southern sub-Saharan Africa have been particularly hard hit, experiencing an explosive late-20th-century spread and subsequent sustained high prevalence. Despite advances in prevention and treatment,

Computational resources used for this work were provided by the American University High Performance Computing System, which was funded in part by the NSF (BCS-1039497). See http://www.american.edu/hpc for information on the system and its uses.

✉ Alan G. Isaac
aisaac@american.edu

[1] Department of Economics, American University, Washington, DC 20016, USA

some regions continue to show startlingly high HIV prevalence (Kharsany et al. 2015). Social scientists have struggled to understand the causes of this epidemic, hoping for insights that can lead to improved health policy. Economists have been interested in the behavioral correlates and the economic consequences of the epidemic (Young 2005; Jones and Klenow 2016).

Efforts to understand the spread of a disease that can be transmitted sexually naturally include investigations of the sexual practices of the affected population. In the 1990s, explanations of the HIV epidemic in sub-Saharan Africa (SSA) began to emphasize the "concurrency hypothesis," which attributes SSA's high prevalence of HIV infection to its high prevalence of long-term concurrent sexual partnerships. The core idea behind this attribution is that overlapping partnerships produce sexual networks that are especially effective at spreading HIV.

Early discussions of concurrency include Watts and May (1992) and Hudson (1993). Advocates of the concurrency hypothesis include Epstein (2007, 2008, 2010), Epstein and Morris (2011), Halperin and Epstein (2007), Mah and Halperin (2010), Morris et al. (2010), and McCreesh et al. (2012). Critics include Lurie and Rosenthal (2010a, b), Sawers and Stillwaggon (2010), Sawers et al. (2011), Tanser et al. (2011), Sawers (2013), and Sawers and Isaac (2017). At present, policy discourse and HIV prevention strategies reflect a belief in the validity of the concurrency hypothesis. For example, Epstein (2008) asserts that " 'long term concurrency' probably explains why HIV in Africa has spread so rapidly beyond typical 'high risk groups' such as sex workers". This view has been driven by SSA's unusual prevalence of HIV and the regional acceptance of polygyny, along with computational models that appear to support an important role for concurrency. However, in the available survey data, correlations between concurrency and HIV prevalence remain empirically fugitive, as most famously illustrated by the "four cities" study of Caraël et al. (2004).[1]

A key strand of the concurrency research constructs agent-based simulation models. Early models demonstrated that concurrency correlates with epidemic scale. Against this, concurrency skeptics have demonstrated that empirically plausible modifications of these models diminish or eliminate the correlations. Critics argue that the concurrency hypothesis has achieved a salience in policy discussions and prevention strategies that is not justified by the empirical and theoretical research. In this article, we join with the critics to raise additional questions about the adequacy of the theoretical work that has been cited in support of the concurrency hypothesis. In order to facilitate comparability of results, our contribution draws on both the structure and the parameterization of some well-known agent-based models of HIV prevalence.

1.1 Agent-based models of the epidemic

Since the pioneering work of Morris and Kretzschmar (1997), agent-based modeling and simulation has come to dominate the theoretical work on the epidemiological consequences of concurrency in sub-Saharan Africa. These models are often called *individual-based pair-formation* models, because they explicitly model individuals

[1] For a detailed discussion, see Sawers and Stillwaggon (2010).

who form and dissolve sexual partnerships. These partnerships are the links across which sexually transmitted disease propagates.

As Eaton (2013, p. 33) puts it, such agent-based models help us understand "how biological and behavioral factors interact in promoting the spread and control of HIV." The Morris and Kretzschmar model—hereafter, the MK1997 model—is a classic demonstration of how greater concurrency prevalence can promote more rapid and effective propagation of sexually transmitted disease. Partly because this linkage is intuitively appealing, their work had a large influence on beliefs about the SSA HIV epidemic.

The MK1997 model is iconic in the literature due both to its methods and to it results. It demonstrated that agent-based modeling can shed light on potential contributors to the SSA HIV epidemic, and it flagged concurrency as a potential culprit. Simulations generated by the MK1997 model are characterized by explosive growth in HIV prevalence, even at low concurrency prevalence. As a result, this model is widely cited in support of the concurrency hypothesis. It has been cited in over 1600 scholarly works on the subject, including subsequent simulation research. Oddly, it took a decade before the model's parameterization began to receive serious critical appraisal (Deuchert and Brody 2007; Lurie and Rosenthal 2010a). During that decade, the concurrency hypothesis became a conventional explanation for SSA's devastating HIV epidemics.

Key aspects of the MK1997 model recur in the subsequent literature. In particular, it is a dynamic, stochastic pair-formation model: partnership formation is ongoing and existing partnerships are constantly at risk of dissolution. Much of the subsequent work on concurrency has followed this general representation of partnership formation and dissolution (Morris and Kretzschmar 2000; Eaton et al. 2011; Sawers et al. 2011; Eaton 2013). We hew very closely to this literature in order to demonstrate how key assumptions have influenced its conclusions. Specifically, we explore the role of partnership duration.

Particularly relevant to the present paper is the work of Morris and Kretzschmar (2000), which extends their earlier MK1997 model. Using the terminology of McCreesh et al. (2012), we can say that this MK2000 model distinguishes between long-duration and short-duration partnerships. Based on survey evidence that long-duration partnerships last about 20 years while short-duration partnerships last about 2 years, the MK2000 model incorporates mean partnership durations that are substantially longer than in the MK1997 model. While HIV epidemics simulated with this model grow much more slowly than in the MK1997 model, prevalence still triples in 5 years (from 1.0 to 2.92%).

The literature contains many modifications of the MK1997 and MK2000 models. The model of Eaton et al. (2011) particularly influences our research.[2] Most importantly, the EHG2011 model incorporates evidence-based transmission rates, which

[2] From the EHG2011 model, we incorporate the improved transmission parameterization, vital dynamics (discussed below), and longer simulation period. At the level of code, we adopt a discrete time-to-event implementation that was inspired by their publicly available C++ code. Our Python implementation is available upon request.

are much smaller than in MK1997.[3] Our model parameterization matches EHG2011 whenever possible, adopting their transmission rates, their vital dynamics (i.e., deaths from AIDS, and the introduction of uninfected individuals), and the substantially longer time-scale of their simulations (chosen to produce simulation results representative of the model's stochastic steady state). However, influenced by the MK2000 model, we additionally distinguish partnership types by duration.

A core modeling goal of the present paper is to maintain results comparability. Our implementation and parameterization therefore closely track existing models (particularly the EHG2011 model). We want the resulting model to be encompassing, so that previous results can be reproduced simply by means of a few parameter changes. In particular, we explore the partnership duration assumptions of the previous literature. To this end, we distinguish primary and secondary partnerships and independently vary the mean partnership duration for each partnership type.[4]

1.2 Exclusivity and partnership duration

It is widely recognized that stable monogamy protects susceptible individuals and traps sexually transmitted diseases within infected partnerships (Dietz 1988; Kretzschmar and Heijne 2017). The same is true of stable polygamy, but many authors have elided the distinction between concurrency and exclusivity. For example, Halperin and Epstein (2004, p. 5) argue that in contrast to a "network of concurrent relationships … serial monogamy traps the virus within a single relationship for months or years". Both trapping of the virus and insulation from it are determined by long-term stability and exclusivity, not by the number of partners, nor by whether the partnerships are formal or informal.

While a stable and exclusive n-person group shares the trapping and protectivity features of a stable and exclusive 2-person partnership, dalliances pose a greater risk for a larger group. Suppose each partner has an independent probability p of an external sexual encounter in a given year. With n partners, the probability that the group remains completely exclusive is $(1 - p)^n$, which is strictly decreasing in the number of partners. Of course the actual risk depends on HIV prevalence, and it is muted by the low per-act transmission efficiency of the virus.

Partnership duration also has important effects on disease spread. We show below that even in the context of concurrent partnering, long partnership duration retards rather than promotes the spread of HIV. Since partnership duration data are scanty, our demonstration relies on examining a range of durations.

Despite numerous surveys of sexual behavior in sub-Saharan Africa, survey evidence on sexual-partnership duration remains scarce. Even surveys that report duration often provide little guidance to modelers. For example, the survey of Powers et al. (2011) is restricted to patients in an STI clinic, who are not representative of the general population. Similarly, surveys restricted to youth will document partnering

[3] They are also appropriately staged (i.e., they vary during the course of the infection.) As explained by Eaton (2013), "[w]ith constant infectiousness, the overall spread of the virus is less dependent on onward transmission within the first few months of infection."

[4] There is only one primary partner, while there may be multiple secondary partners.

patterns unlikely to generalize to an older population (Harrison et al. 2008; Harrison and O'Sullivan 2010; Goodreau et al. 2012). Studies of youth also produce right censoring of partnership duration.

However, recent Demographic and Health Surveys include some useful questions about primary partnership duration. In 16 countries in sub-Saharan Africa, women reported mean duration in married or cohabiting partnerships of 12.1 years, and men reported mean duration of 13.1 years. This is virtually the same as in 16 low-income countries outside Africa for which there are data (based on datasets obtained from The DHS Program, ICF International).

Unfortunately, data on secondary partnerships are not reported by the DHS. There has been some attempt to address this in in the literature on formal polygyny (Reniers and Tfaily 2008, 2012; Reniers and Watkins 2010). Of particular interest are the results of a survey in Rakai, Uganda, as reported by Morris and Kretzschmar (2000, Table 1). To our knowledge, this the only survey that reports average duration of both primary and secondary partnerships for a large sample of adults ($N = 1994$), and it found that the average duration of secondary partnerships was about 12% of average duration of primary partnerships (28.4 months vs. 239.1 months). This difference is roughly an order of magnitude, which suggests that models of sexually transmitted diseases should explore the implications of varying partnership duration by partnership type.

The available survey evidence is even less informative about the interaction of duration, concurrency, and HIV. Our model sheds light on these interactions, which may suggest directions for future data collection and empirical design. We thereby contribute to the agent-based simulation and modeling literature that elucidates how partnership characteristics influence the spread of sexually transmitted disease (Chen et al. 2008; Kim et al. 2010).

2 Model overview

This section describes our agent-based two-sex pair-formation model of the heterosexual propagation of HIV. To facilitate comparability with previous research results, the model adheres closely to the EHG2011 model surveyed in the previous section and to the SIS2011 model discussed below.

2.1 Pair formation in an agent-based model

As in the MK1997 and EGH2011 models, we consider a 2-sex population (males and females). Figure 1 provides a simplified schematic, where each `Person` has an immutable `sex` attribute.[5] Each person has a `partnerships` list, which represents the current set of sexual partners. Under monogamy, an individual has at most one

[5] This heuristic schematic is essentially a simplified UML class box (Bersini 2012), but that is not meant to imply any particular approach to implementation. It documents only key public features, specifies only the types of operation (e.g., method) arguments, and omits constructors. Conventionally a `frozen` property (in braces) indicates immutability, but this is omitted to reduce diagram clutter. Operations specify no return type because they are called only for their side effects.

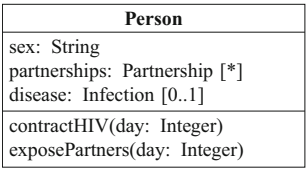

Fig. 1 Three main attributes and two main operations of a `Person`

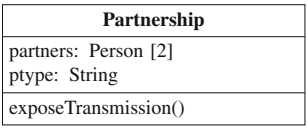

Fig. 2 Two main attributes and one main operation of a `Partnership`

partnership; under concurrency, an individual may have multiple concurrent partnerships.[6]

The model employs a natural implementation of disease transmission: on any given day, an individual can contract a disease via the `contractHIV` method. We are modeling the transmission of HIV, which is permanent (until death), so an individual contracts it only once.[7] When an individual contracts the disease, all of the individual's partners are exposed (via `exposePartners`), which may lead to disease transmission. In pair-formation models, it is perfectly feasible to track the force of infection by attending to the stages of a disease and the frequency of sex. Eaton et al. (2011) and Sawers et al. (2011) demonstrate that these considerations can be important for epidemic spread, and our model incorporates their insights.

In a *serodiscordant* sexual partnership, an infected individual exposes a susceptible individual to infection. This paper models the sexual transmission of HIV at the individual level, introducing explicit links between the individuals. Such *pair-formation* models, were an important innovation in infectious disease modeling in the 1980s (Kretzschmar and Heijne 2017). Unlike classical epidemiological models, pair-formation models can accommodate repeated sexual contact within long-duration partnerships of varying degrees of exclusivity, which has proved important for understanding the dynamics of sexually transmitted infections (Dietz 1988; Kretzschmar and Dietz 1998). The HIV epidemic in sub-Saharan Africa (SSA) appears to be primarily spread by heterosexual sex, so as in the MK1997 and EHG2011 models, we consider only heterosexual partnerships.[8]

Figure 2 provides a simplified `Partnership` schematic, representing a heterosexual sexual relationship. A partnership most essentially is a pair of individuals.[9] However, sexual partnerships may be of different types, which we capture with a

[6] The `[*]` annotation conventionally indicates that there can be zero or more partnerships. (We use an ordered collection to ensure replicability.) Partnering is dynamic: an individual may add or remove a partnership. (To avoid clutter, we do not list the associated operations.) Partnership formation and dissolution is stochastic, as specified below.

[7] This simplifying assumption is standard in HIV modeling. In fact there are two main types of HIV (and numerous strains), and individuals can be infected by more than one type or strain.

[8] Partnerships are therefore the undirected edges in a dynamic bipartite graph.

[9] The fixed multiplicity of 2 is indicated in brackets, following a UML convention. The pair of `partners` and the `ptype` are immutable attributes of a partnership.

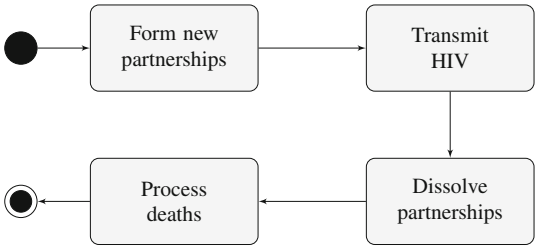

Fig. 3 Daily simulation schedule: new partnerships form, HIV is transmitted, existing partnerships dissolve, and some individuals die from HIV infections

`ptype` attribute. Morris and Kretzschmar (2000) included *long-term* and *short-term* partnerships; Kretzschmar (1995) and Bauch (2002) distinguish *steady* and *casual* partnerships; and Fu et al. (2016) distinguish *regular* and *casual* partnerships. We adopt the terminology and approach of Sawers et al. (2011); their SIS2011 model distinguishes *primary* and *secondary* partnerships.

Primary partnerships resemble the regular partnerships of Fu et al. (2016), in that individuals can have at most one primary partnership at a time, while an individual may have multiple secondary partnerships. There is no implication that a primary partnership is somehow institutionally sanctioned or socially recognized, that primary partners cohabit, or even that primary partnerships last longer than secondary partnerships. In principle, primary partnerships may differ from secondary partnerships along many dimensions. For example, Sawers et al. (2011) considered differences in the frequency of sex, and Sawers and Isaac (2017) additionally allow for a proportional difference in mean partnership duration. The present paper focuses of the role of partnership duration, which is allowed to differ arbitrarily between primary and secondary partnerships. (Coital dilution is separately considered as an extension in Sect. 3.2.)

2.2 Model dynamics

Our core simulation schedule is standard. Figure 3 illustrates the schedule for a single day. Each day, individuals stochastically form partnerships. Serodiscordant partnerships probabilistically transmit HIV. Partnerships may dissolve, and infected HIV individuals may die from their disease. Each component of this schedule incorporates stochastic elements.

The literature contains a variety of approaches to partnership formation, but all conceptually involve random encounters between individuals who prefer not to be single. Each period, any individual may encounter a possible partner, which may lead to the formation of a partnership. Given an opportunity, two single individuals will form a pair. However, partnership formation may not be ensured if partnering will produce concurrency. For example, two monogamous individuals will not form a partnership if either already has a partner.

Table 1 Partnership formation

First person's number of partners	0	0	1+	1+
Second person's number of partners	0	1+	0	1+
Partnership formation probability (ϕ)	1	$1-\varepsilon$	$1-\varepsilon$	$1-\varepsilon$

Conditional on an opportunity, partnership-formation probability depends on the level of concurrency resistance ($\varepsilon \in [0..1]$)

Our model is part of the concurrency literature, where partnership formation is effectively the outcome of a negotiation that is subject to random influences.[10] Following the MK1997 model and its offshoots, concurrency resistance (ε) parameterizes the extent to which concurrency is an obstacle to partnership formation. As in the EHG2011 model and SIS2011 model, this paper adopts the random-mixing approach of Kretzschmar and Morris (1996, p. 180) summarized in Table 1. The probability of partnership formation (conditional on an opportunity) depends on the current number of partners and social norms, which are captured by the level of concurrency resistance. As shown by Eaton et al. (2011), mean concurrency outcomes depend on the value of ε. One may therefore vary ε in order explore the relationship between concurrency and epidemic outcomes.

If $\varepsilon = 1$, concurrency is fully resisted, and only monogamous partnerships can form. If $\varepsilon = 0$, concurrency is not resisted; existing partnerships do not influence the formation of new partnerships.[11] Equation (1) gives the decision table an algebraic expression matching Kretzschmar and Morris (1996, eq. 40). The partnership formation probability, ϕ, is parameterized by the resistance-to-concurrency parameter, $\varepsilon \in [0 \ldots 1]$. Here $\deg(p_i)$ denotes the number of partners of person p_i.

$$\phi(p_1, p_2, \varepsilon) = \begin{cases} 1 & \deg(p_1) = 0 \text{ and } \deg(p_2) = 0 \\ 1-\varepsilon & \text{otherwise} \end{cases} \tag{1}$$

Allowing for concurrency dramatically increases the maximum number of partnerships that are logically possible at any time. This introduces some subtle modeling considerations. Kretzschmar and Morris (1996) argue that in order to isolate the contribution to HIV spread of concurrency per se, one must limit the maximum number of pairs under concurrency to the number that is possible in the model under monogamy. This is because tying concurrency to an increase in the *society-wide* number of partnerships

[10] One is reminded of a description in Binmore (2007, p. 536): "One may imagine that Alice knocks on doors at random if unmatched. When Bob or Chris answers the door, she bargains with him until agreement is reached or their unpredictable wives grow tired and run her off the property."

[11] Following Eaton et al. (2011), we do not impose a maximum number of partners for each individual. Other approaches are common in the literature. For example, Morris and Kretzschmar (1997) impose a maximum of 4 partners per individual. Morris and Kretzschmar (2000) report Ugandan survey data that suggests that having even as many as 3 partners is rare, and their associated simulation model correspondingly imposes a maximum of 3 partners. Caraël et al. (2004, p. 65) report that more than 8 partners in a year is very rare, with nearly half of respondents reporting 0 partners and almost all the rest 5 or fewer. Although we do not impose an exogenous maximum, individuals with more than a few partners are extremely rare, even when there is no concurrency resistance whatsoever.

would automatically increase the likelihood of transmission, obscuring the specific role of concurrency. They argue that investigations of the role of concurrency should therefore explore the role of the distribution of partnerships separately from the total number of partnerships. Following MK1997 and EHG2011, we accept this argument.[12]

Each serodiscordant partnership risks disease transmission. The magnitude of transmission risk is therefore a crucial determinant of epidemic outcomes. As a famous example, Morris and Kretzschmar (1997) demonstrate that an unrealistically high 0.05 daily infection risk will produce an explosive HIV epidemic. Therefore, following Eaton et al. (2011), we incorporate available evidence on the time-dependent daily risk of HIV infection.[13]

The end of a partnership ends the risk of disease transmission in that partnership. Our model of partnership dissolution follows the standard practice in the literature: each period there is a risk that a partnership will not survive. This means that the duration of each partnership is idiosyncratic: some are high duration, and some are brief duration. In this paper, we are particularly interested in the influence of partnership duration on epidemic spread. In order to better assess the claim that long-duration concurrent partnering promotes HIV spread, we will consider substantial variations in mean partnership duration, and we will also allow duration to vary by partnership type.

The final action in the daily simulation schedule is to remove from the simulation the individuals who die of HIV. We follow the common practice in this literature and maintain a constant population: an agent who dies of HIV is replaced by a new uninfected agent of the same sex.[14]

[12] See the "Appendix" for a detailed parameterization. See Kretzschmar and Morris (1996, p. 181) and Eaton (2013, p. 48) for a detailed discussion and corresponding pseudocode, or Eaton et al. (2011) for a complete C++ implementation. This approach is very common in the literature, but variations exist. For example, the MK2000 model instead sets a ceiling on total partnerships equal to the maximum that is possible given their ceiling on the number of partnerships per person. However, that model is applied to a different question (the consequences of the mean degree of the population in various scenarios).

[13] In an agent-based model, one might expect to work with a per-act risk of infection, possibly deriving a daily risk of infection as the product of a per-act transmission rate and coital frequency. However, Hollingsworth et al. (2008) argue that measuring the number of sex acts is subject to substantial error, which produces a corresponding error in the per-act transmission rate. In contrast, there is substantial evidence from seroconversion data concerning annual transmission risks. Working directly with daily risk of infection allows less problematic calibration to this evidence. We therefore adopt this approach, which has remained fairly standard in the literature since the MK1997 model. Specifically, we follow the EHG2011 model in calibrating the daily transmission risk to the stage-specific annual risk of infection documented by Hollingsworth et al. (2008).

[14] Following Eaton et al. (2011) and Sawers et al. (2011), our individuals do not die of any other causes. (Contrast with McCreesh et al. (2012), who remove individuals from the model at age 55.) Allowing individuals to die of old age is largely irrelevant for our research questions, although it would very slightly reduce transmission likelihood (since in principle an infected partner could die of old age before transmitting the disease).

Additionally, as noted by two referees, our baseline model does not include mother-to-child transmission. In a scenario contrast experiment (not shown), we allow seropositive mothers to produce seropositive offspring who survive to adolescence. In comparison with our baseline results, the duration results with seropositive births are qualitatively similar. As expected, this change boosts prevalence at every duration combination.

3 Parameterization and results

This section discusses the model parameterization and lays out the focal results. It then examines a set of scenario-contrast experiments, in order to demonstrate the robustness of these results.

3.1 Baseline

In order to maximize results comparability, our baseline parameterization matches Eaton et al. (2011, Table 1) extremely closely. We include three deviations from their choices: the number of agents, the range in the levels of concurrency resistance, and most importantly, the treatment of partnership duration. (See the "Appendix" for more parameterization details.)

The crucial change is the treatment of partnership duration, which is the focus of our paper. The MK1997 and EHG2011 models impose a mean partnership duration of 200 days. As discussed in our literature review, when compared to the partnership durations in SSA data, this proves implausibly brief (Morris and Kretzschmar 2000). Additionally, we wish to review the claim that long-duration concurrency is key to spreading HIV. We therefore allow the mean partnership duration to vary from 1 to 20 years. To facilitate comparison with previous literature, we additionally consider a duration of 200 days (Kretzschmar and Morris 1996; Eaton et al. 2011; Morris and Kretzschmar 2000; Sawers et al. 2011). Additionally, we distinguish between primary and secondary partnerships, as discussed in the previous section, and we allow each partnership type to have a different mean duration.[15]

In a second deviation from the EHG2011 parameterization, our focal experiment does not vary concurrency resistance. Whereas Eaton et al. (2011) consider 11 different values for concurrency resistance, our focal simulations consider just one. This is because the EHG2011 results already demonstrate that epidemics fail to spread at high levels of concurrency resistance. Our focal simulations therefore avoid reconsidering this issue, and we consider only the lowest level of concurrency resistance in the EHG2011 scenarios (which correspondingly produces the highest concurrency prevalence). Whenever we find that epidemics fail to emerge, our results are correspondingly strong. (Sect. 3.2 nevertheless reports results for some additional values of concurrency resistance, as scenario-contrast experiments.)[16]

Finally, our focal simulations use one-tenth the population of the EHG2011 simulations. This is meant to enhance the replicability of our core results by reducing the

[15] These models determine mean duration by drawing each partnership's duration from a (shifted) geometric distribution with parameter $\sigma = 0.005$, the daily risk of dissolution. In contrast, we will consider 21 different values of σ for each partnership type, allowing that primary partnerships may be longer or shorter in duration than secondary partnerships. (Our distinction between primary and secondary partnerships therefore does not correspond to the distinction between long-duration and short-duration partnerships in Morris and Kretzschmar (2000) or McCreesh et al. (2012), although it encompasses it.)

[16] Results for the full EHG2011 range of concurrency resistances are available upon request.

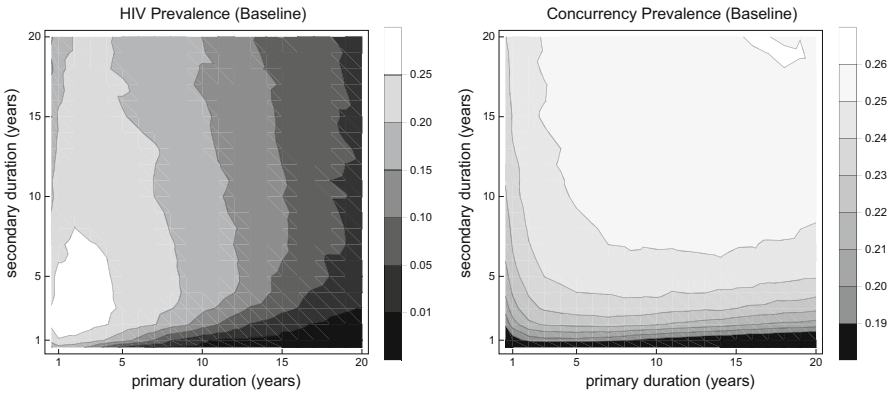

Fig. 4 Average outcomes for baseline parameterization: final values of HIV prevalence and concurrency prevalence depend on the mean duration of primary partnerships and of secondary partnerships. Prevalence levels are shaded between contours, as indicated by the adjacent legends

(still substantial) required simulation time.[17] However, to demonstrate robustness to scale, Sect. 3.2 additionally provides scenario-contrast experiments for variations in the number of agents.

As in Eaton et al. (2011) and Sawers et al. (2011), we produce an initial HIV infection by randomly infecting 1% of the population. The entire simulation schedule then runs daily for 250 years. (We additionally consider shorter time horizons in Sect. 3.2.) Figure 4 reports end-of-simulation mean HIV prevalence and mean concurrency prevalence for the focal experiment. (Means are across 100 replicates, so the figure displays results for the 44, 100 simulations run with the baseline parameterization.) The first subplot is a contour plot of the final HIV prevalence for each considered duration combination (primary-partnership mean duration and secondary-partnership mean duration, measured in years). Some nonmonotonicity is evident, but these plots nevertheless readily suggest some generalizations.

Most obviously, an increase in the mean duration of primary-partnerships tends to reduce the likelihood of epidemic spread. Furthermore, if we select an arbitrary point in the plot and then simultaneously increase primary-partnership duration and secondary-partnership duration, HIV prevalence generally declines. Similarly, we usually see a decline even if we increase only the secondary-partnership duration. However, if we pick an initial point that has low secondary-partnership duration and then increase it only modestly, the result may be an increase in HIV prevalence. This is understandable: as noted in Sawers and Isaac (2017), in a partnership with a very brief duration, HIV transmission is quite unlikely. Nevertheless, as a rough generalization, Fig. 4 shows both for primary partnerships and for secondary partnerships that increased partnership duration constrains epidemic spread. Furthermore, the figure encompasses both the duration values used by MK1997 and EHG2011 and the empirically grounded

[17] A single experiment comprises 441 distinct scenarios (i.e., combinations of primary-partnership mean duration and secondary-partnership mean duration). Our focal experiment includes 100 replicates of each scenario, which—for this focal experiment alone—implies more than 40, 000 simulations of *daily* actions over our time horizon of 250 years.

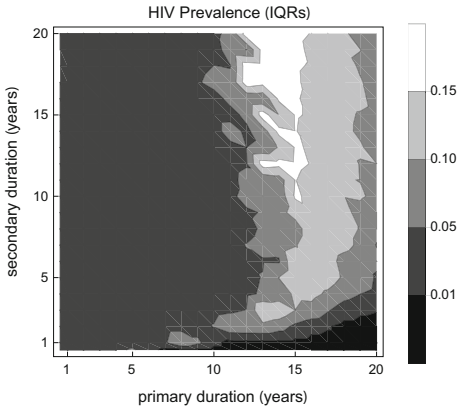

Fig. 5 Uncertainty in baseline results: IQRs for HIV prevalence vary with the mean duration of primary partnerships and of secondary partnerships. An IQR across replicates is computed for each duration scenario, and the results are shaded between contours (as indicated by the adjacent legend)

partnership durations suggested by Morris and Kretzschmar (2000). The former produce epidemic spread, while the latter produce epidemic extinction. This difference highlights the crucial role of duration assumptions for the results in the existing literature.

The second subplot of Fig. 4 offers a look at concurrency prevalence, allowing us to ensure that our epidemic results do not reflect some unexpected response of concurrency prevalence to changes in partnership duration.[18] This figure shows that concurrency prevalence is in fact extremely high in this population at every combination of partnership durations. (Recall that for these results we used the *lowest* level of concurrency resistance considered in EHG2011; for more plausible levels of concurrency, see the scenario-contrast experiments below.) Unsurprisingly, concurrency prevalence tends to rise as secondary-partnership duration increases, since it becomes more likely that someone with a partner acquires an overlapping partner. This means that, as seen by comparing the two subplots in Fig. 4, concurrency is actually highest where HIV prevalence is lowest. In striking contrast to the MK1997 and EHG2011 conclusions, our focal results provide no reason for empirical researchers to expect a positive simple correlation between concurrency prevalence and HIV prevalence.

The occasionally rough contour edges in the Fig. 4 hint at some underlying variability in the reported simulation outcomes. To better characterize this variability, we adopt a popular robust measure: the interquartile range (IQR), which is the difference between the third and first quartiles of a sorted data set. For each scenario represented in the figure, we compute the IQR across replicates.

Figure 5 presents these IQRs for HIV prevalence under the baseline parameterization. These are most often low. However, when both primary partnerships and secondary partnerships endure approximately 15 years, we find evidence of a transition region, where where epidemics and extinctions are both frequent. In the regions typical of the literature and in the most empirically plausible regions, IQRs are small.

[18] Recall that EHG2011 has a single partnership duration for all partnerships, which is an unrealistically short 200 days. This constrains the possible levels of concurrency prevalence. With more plausible (much longer) partnership durations, we naturally see higher concurrency prevalence.

We do not show corresponding results for concurrency prevalence, since the IQRs are uniformly low (less than 0.01).

3.2 Exploring the model

This subsection presents a collection of scenario-contrast experiments that shed light on the scope and robustness of our reported results. Additionally, as a preliminary exploration, we briefly examine the emergence of HIV epidemics in the baseline model.

3.2.1 The first 100 years

Eaton et al. (2011) reported results after 250 years, in order to ensure they were adequately representing the stochastic steady state of the simulation. The present paper follows this practice for the same reason. An important additional reason is to facilitate comparability with the EHG2011 and SIS2011 results. Nevertheless, the emergence of the epidemic is also of interest. For example, the speed of emergence is important for assessing the immediacy of any public health crisis. Additionally, it may be tempting to dismiss outcomes in the distant future on the basis of ongoing efforts to reduce transmission rates with ART and PrEP or of presumed advances in medical technology.[19] Figure 6 therefore illustrates the emergence of the epidemic at 25 year intervals over the first 100 years.[20]

We find that the ultimate pattern of HIV prevalence (by partnership duration) emerges almost completely in the first 100 years of the epidemic. It is instructive that the most discouraging outcomes after 25 years are only a harbinger of worse things to come. At 25 years, the worst outcomes are achieved when all partnerships are only a few years in length, and this region continues to experience the worst epidemic spread over time.

3.2.2 Population size

By considering substantially smaller and substantially larger populations, this subsection demonstrates that our results are robust to population size. Our baseline parameterization has 2000 individuals. Here we present results for 400 individuals and 20, 000 individuals. Aside from the change in population size, we maintain our baseline parameterization.

First consider the results for the larger population. Unsurprisingly, a larger population brings substantially decreased variability in outcomes. We therefore report results for 50 replicates, in order to facilitate replicability. (Similar results with 100 replicates are available upon request.) The first subplot of Fig. 7 shows the end-of-simulation HIV prevalence outcomes for this larger population. Aside from the much smoother contours, the results are very similar to our focal simulations, reported in Fig. 4.

[19] For example, the Partners Demonstration Project found that integrated ART and PrEP delivery to serodiscordant couple drastically lowered HIV transmission (Baeten et al. 2016).

[20] We do not illustrate concurrency prevalence since (as expected) the outcomes are not changed.

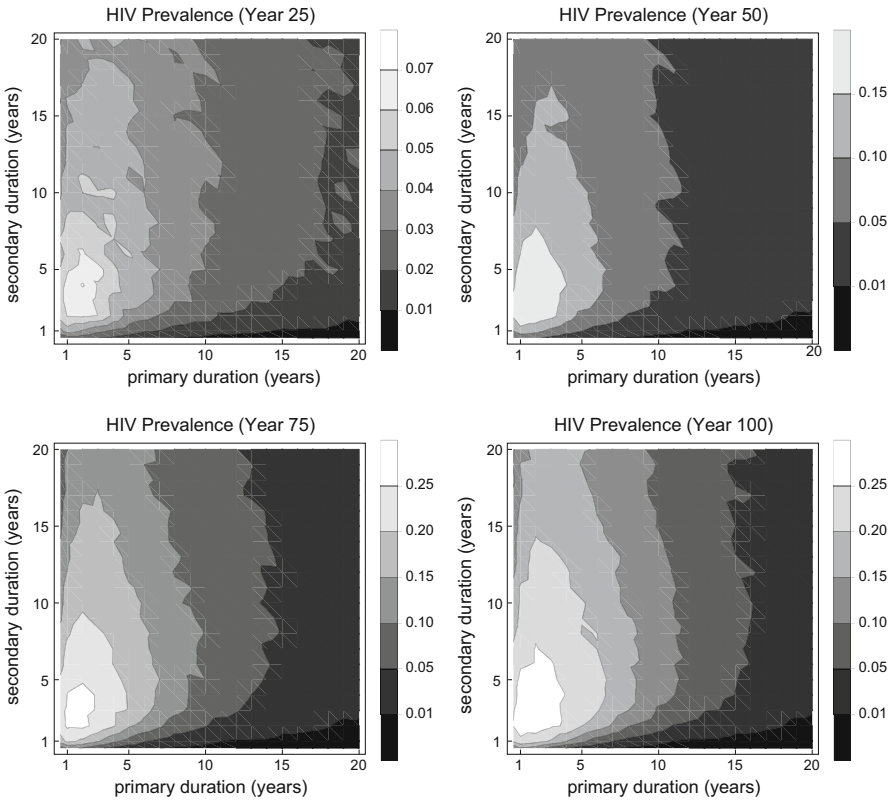

Fig. 6 Epidemic emergence (baseline parametrization): HIV prevalence grows during the first 100 years, particularly with low duration partnerships

Once again there is no support for claims that long-duration concurrency contributes to epidemic spread. A smaller population brings substantially increased variability in outcomes, so we increased the number of replicates to 500. The second subplot of Fig. 7 shows the end-of-simulation HIV prevalence outcomes with this diminished population. While it too resembles Fig. 4, epidemics clearly have more trouble getting started in this smaller population. Nevertheless, as in our focal results, the HIV-prevalence outcomes respond negatively to increases in partnership duration.

Concurrency outcomes change little in either scenario, so we do not report them. The means that, as in our focal results, the lowest levels of mean HIV-prevalence correspond to to the highest levels of concurrency in the population.

3.2.3 Coital dilution

Our model extends the EHG2011 model to explore the implications of partnership duration. Sawers et al. (2011) extend this same model to argue that coital dilution can have an important effect on HIV transmission. This subsection therefore briefly addresses the importance of coital dilution.

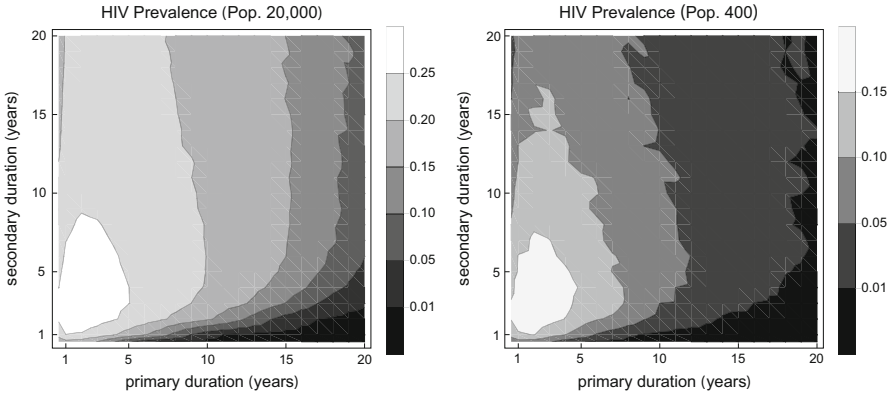

Fig. 7 Robustness to population size: in populations that are larger or smaller than the baseline, HIV prevalence similarly depends on partnership durations

Coital dilution is a reduction in the frequency of sex per partner for individuals with multiple partners. While data on coital frequency in sub-Saharan Africa and elsewhere are thin, Sawers et al. (2011) summarize evidence suggesting substantial coital dilution in the presence of multiple partners. As one example, respondents in the Rakai survey reported coital dilution of more than 75% (Morris et al. 2010).[21] However, some authors contest the existence of coital dilution (Delva et al. 2013; Jenness et al. 2015), and directly allowing for coital dilution is rare in the theoretical work on HIV transmission. We therefore include coital dilution as a scenario-contrast experiment, rather than in our focal simulations.

Sawers et al. (2011) found that coital dilution constrains the spread of HIV and thereby inhibits the emergence of epidemics. Here we illustrate similar epidemic-dampening effects of coital dilution by adopting a rather conservative 25% coital dilution rate for secondary partnerships. As in the SIS2011 model, we introduce coital dilution by reducing the transmission rate in secondary partnerships. Aside from this change, we duplicate the baseline parameterization used in our focal simulations, reported above.[22]

The first subplot of Fig. 8 shows the end-of-simulation HIV prevalence outcomes with coital dilution. In support of the results of Sawers et al. (2011), we find that coital dilution has an important effect on epidemic likelihood in HIV simulations. In comparison to our focal simulations reported in Fig. 4, we see substantial dampening of epidemic outcomes. The SIS2011 result is not overturned by our consideration of variations in partnership duration. It is noteworthy that even with primary-partnership duration of a little as ten years, there is no level of secondary-partnership duration that leads to an epidemic.

[21] For additional evidence of coital dilution, see Gaydosh et al. (2013), Reniers and Tfaily (2010), Reniers and Tfaily (2012), and Reniers and Watkins (2010).

[22] Primary partnerships have the highest frequency of sex but need not have the longest duration. As a concrete example, the post-war partnership between Sartre and de Beauvoir might be classified both as long-duration and as secondary.

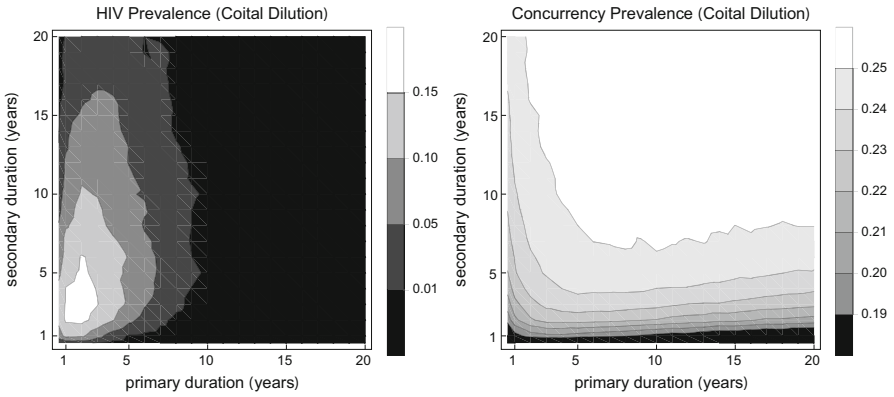

Fig. 8 Effect of coital dilution: relative to the baseline parameterization, coital dilution lowers HIV prevalence at every combination of partnership durations. Concurrency prevalence does not change

The second subplot of Fig. 8 shows the end-of-simulation concurrency prevalence outcomes with coital dilution. Unsurprisingly it closely resembles Fig. 4. Once again, as in our focal results, the lowest HIV-prevalence outcomes correspond to the highest levels of concurrency in the population.

As discussed above, we have so far accepted the Kretzschmar and Morris (1996) admonition to hold constant the maximum number of partnerships in our pair-formation model of concurrent partnering. However, as emphasized by Eaton (2013, p. 60) the introduction of coital dilution begs the question of whether this admonition should apply to the total number of partners or to the total amount of sex in the population. For example, in a model of gonorrhea, Welch et al. (1998, p. 245) introduce a "concurrency adjusted sex-act rate" with the goal of holding constant the total number of sex acts in the population. This approach may be seen as more in tune with the notion of individual "sex budgets" introduced in Blower and Boe (1993). Without taking a position on which approach to coital dilution leads a theoretical model to most helpfully guide empirical explorations, it is clear that the consequences within a model for HIV spread must be quite different. The results reported here emphasize comparability with the SIS2011 model.

3.2.4 Concurrency resistance

The focal simulations reported in Sect. 3.1 incorporate the lowest level of concurrency resistance considered in the EHG2011 and SIS2011 models. The corresponding levels of concurrency prevalence, displayed in the second subplot of Fig. 4, are therefore very high. However, Sawers and Isaac (2017) suggest that realistic rates of concurrency prevalence are 20% or less—quite plausibly substantially less.

This subsection therefore considers the implications of somewhat higher levels of concurrency resistance—enough to push the lowest concurrency outcomes into an empirically plausible region. The results are displayed in Fig. 9. With more plausible levels of concurrency, there is a much smaller region of partnership durations in which an epidemic may emerge. Two key features of our focal simulation remain visible:

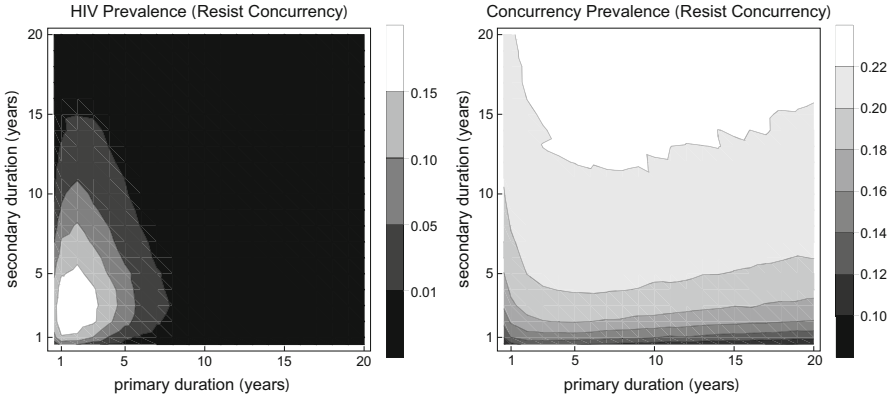

Fig. 9 Effects of concurrency resistance: with greater concurrency resistance, HIV prevalence falls at all partnership duration combinations. Concurrency prevalence also falls

an increase in partnership duration generally reduces epidemic likelihood, and the highest concurrency rates coincide with the lowest epidemic rates.

3.2.5 Concurrency and exclusivity

Our previous results treat the issue of concurrency resistance as in EHG2011 and SIS2011. As a final scenario-contrast experiment, we reformulate concurrency resistance so that it exhibits network sensitivity. As emphasized by Caraël et al. (2004, p. 60) "people are put at risk not just by their own behavior but by that of others to whom they are linked", and (p. 71) "the partners of someone who has several partners may not have other partners." In this experiment, our agents display network sensitivity in partnership selection. Specifically, although agents refuse to partner with those whose partners have concurrent partners, multiple partnering is otherwise not resisted at all.

This network sensitivity implies a modification of our partnership-formation specification. Although there is no direct resistance to concurrency, there is resistance to overlapping webs of concurrent partnerships. For example, under this specification, a woman would not resist joining a (formal or informal) polygynous marriage, as long as none of the other wives have multiple partners. However, individuals refuse to partner with someone whose partners have multiple partners. Table 2 provides a decision-table formulation of the new probability of partnership formation (ϕ^*). Equation (2) provides an equivalent algebraic formulation of the process, where $\deg^*(p)$ denotes the maximum number of current partners of any partners of person p.

$$\phi^*(p_1, p_2, \varepsilon) = \begin{cases} 1 & \deg(p_1) = 0 \text{ and } \deg^*(p_2) < 2 \\ 1 & \deg^*(p_1) < 2 \text{ and } \deg(p_2) = 0 \\ 1 - \varepsilon & \text{otherwise} \end{cases} \tag{2}$$

Table 2 Partnership formation with network sensitivity

d2(p_1)	0	0	0	1	1	1	2	2	2
d2(p_2)	0	1	2	0	1	2	0	1	2
ϕ^*	1	1	$1-\varepsilon$	1	$1-\varepsilon$	$1-\varepsilon$	$1-\varepsilon$	$1-\varepsilon$	$1-\varepsilon$

Conditional on an opportunity, partnership formation likelihood depends on the level of concurrency resistance ($\varepsilon \in [0..1]$) and the structure of partnerships. Definitions: d2 = 0 means no partners; d2 = 1 means one or more partners, but *those* partners have no other partners; d2 = 2 means at least one partner has another partner

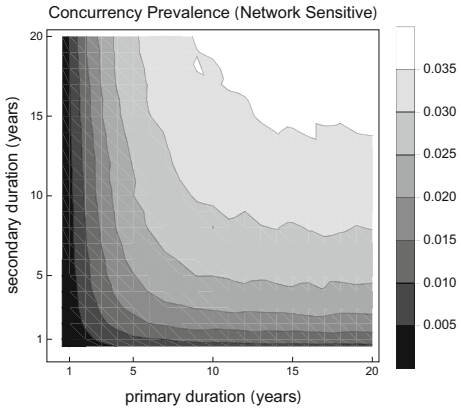

Fig. 10 Effects of network sensitivity: concurrency prevalence falls dramatically (relative to the baseline) when partnership formation is network sensitive

We do not show a contour plot for the end-of-simulation HIV prevalence outcomes with network-sensitive concurrency resistance, because no epidemics arise at any combination of partnership durations. A key reason for this is rather surprising: concurrency plummets with this partnership-formation specification. Figure 10 shows the end-of-simulation concurrency prevalence outcomes with network-sensitive concurrency resistance. Despite the complete absence of concurrency resistance when it comes to star networks, concurrency prevalence is extremely low at every partnership-duration combination. This is because whenever an agent's partner has another partner, the agent is no longer attractive for partnership formation.

4 Conclusion

Assumptions about partnership duration have a large influence on the results of agent-based pair-formation models of concurrency and HIV spread. We demonstrate the size of this influence by adapting the well-known EHG2011 epidemiological model (Eaton et al. 2011). Even at the lowest level of concurrency resistance considered in this model—which produces extremely high levels of concurrency—partnership durations determine whether or not an initial HIV infection will result in an epidemic.

Our results contradict claims that long-term concurrent partnering is an important contributor to HIV spread. In general, a longer mean partnership duration slows the pace at which simulated epidemics grow. We allow mean partnership duration to vary between primary and secondary partnerships. With empirically plausible assumptions

about partnership duration and at levels of concurrency found in sub-Saharan Africa, simulated HIV epidemics grow slowly or not at all.

Our results are consistent with the hypothesis that long-duration partnering is protective against HIV and inconsistent with the hypothesis that long-term concurrency drives the HIV epidemics in sub-Saharan Africa. Additionally, our results provide no reason for empirical researchers to expect a positive simple correlation between concurrency prevalence and HIV prevalence. Nevertheless, holding mean partnership duration constant, an increased reluctance to form concurrent relationships produces a reduction both in concurrency and in the likelihood of epidemic spread.

These results have implications for empirical work on the relationship between concurrency and HIV prevalence. We have seen that even at extremely high concurrency levels, there is no epidemic spread when partnership duration is high. Our results establish the importance for empirical work of recognizing that both HIV prevalence and concurrency prevalence are endogenous outcomes. Behaviors that increase concurrency prevalence—here, average partnership duration—can reduce HIV prevalence. Empirical as well as theoretical studies of concurrency and HIV should therefore control for partnership duration. Our results raise particularly serious questions for researchers whose model parameterizations impose unrealistically short values for mean partnership duration. This would include the MK1997 model and the EHG2011 model, but not the MK2000 model.

We also find that coital dilution has a strong effect on epidemic likelihood. As with the SIS2011 model Sawers et al. (2011), our results suggest that the empirical debate over the extent of coital dilution will have important implications for HIV modeling. Finally, we find that resistance to concurrency is much less important than commitment to exclusivity in limiting HIV spread.

A core question in HIV research is the status of the concurrency hypothesis. Is long-term concurrency a primary explanation of HIV's rapid 20th century growth and persistently high HIV prevalence in certain countries of sub-Saharan Africa? Despite closely aligning our model of HIV transmission and concurrent partnering with key models in the literature that have supported the concurrency hypothesis, our results raise serious doubts about this hypothesis. In particular, we show that long duration partnering is protective, even in the presence of unrealistically high levels of concurrency. We also demonstrate that, in the presence of concurrency, longer duration partnering implies more concurrency but less HIV. In this sense, our findings suggest that the two-decade effort to explain HIV prevalence in sub-Saharan Africa by focusing on the role of concurrency has been misguided. Instead, we lend support to a substantial literature asserting that stable polygyny can be protective against HIV (Blower and Boe 1993; Kretzschmar et al. 2010; Lurie and Rosenthal 2010a; Reniers and Tfaily 2012; Reniers and Watkins 2010).

A detailed parameterization

This appendix provides a detailed parameterization for the baseline model. References to EHG2011 are to Eaton et al. (2011). References to MK1997 are to Morris and Kretzschmar (1997). References to HAF2008 are to Hollingsworth et al. (2008).

Note: EHG2011 effectively has only primary partnerships, so primary and secondary partnership parameters are identical in the baseline parameterization below.

```
[nMale]
type : int
value : 1000
source : see paper
description : number of males in the population
[nFemale]
type : int
value : 1000
source : see paper
description : number of females in the population
[nRandomSeed]
type : int
value : 7930881
source : EHG2011's R code
description : base random seed for each scenario; incremented per replicate
[rho]
type : float
value : 0.01
source : MK1997, via EHG2011 Table 1
description : partnership formation rate
[pSeedHIV]
type : tuple of float
value : 0.01 0.01
source : EHG2011 Table 1
description : proportion of males and females seeded with HIV
[nBurnDays]
type : int
value : 365*5
source : EHG2011
description : number of burn days before disease seeding
[nSimDays]
type : int
value : 365*250
source : EHG2011
description : number of simulation days after seeding
[transmissionRatesHIV]
type : tuple of float
value : 0.007315068 0.0002904110 0.002082192 0.0
source : the file concprimaryinf/R/conc.sim.R accompanying EHG2011
description : average (m/f) DAILY transmission probabilities during each stage
[durationsHIV]
type : tuple of int
value : 88 3054 274 307
source : HAF2008 via EHG2011, Table 1
description : duration (in DAYS) of infection stages
[duration01]
type : float
value : 200
source : MK1997, via EHG2011 Table 1 (equates to sigma of 0.005)
description : mean duration (in DAYS) of primary partnerships
[duration02]
type : float
value : 200
source : to match EHG2011 Table 1 (which does not differentiate pship types)
description : mean duration (in days) of secondary partnerships
[secondarySexfreq]
type : float
value : 1.0
source : chosen to match EHG2011 (which does not differentiate pship types)
description : relative sex-act frequency (secondary vs primary pships)
[epsilon]
type : computed float
value : None
source : MK1997, via EHG2011 Table 1, float in (0,1), see paper
description : concurrency resistance parameter
```

```
[pTransmissionMotherChildHIV]
type : float
value : 0.0
source : EHG2011 has no MTCT, so 0.0 baseline probability
description : probability of seropositive mother-to-newborn transmission
```

References

Baeten JM, Heffron R, Kidoguchi L, Mugo NR, Katabira E, Bukusi EA, Asiimwe S, Haberer JE, Morton J, Ngure K, Bulya N, Odoyo J, Tindimwebwa E, Hendrix C, Marzinke MA, Ware NC, Wyatt MA, Morrison S, Haugen H, Mujugira A, Donnell D, Celum C, Team PDP (2016) Integrated delivery of antiretroviral treatment and pre-exposure prophylaxis to HIV-1-serodiscordant couples: a prospective implementation study in Kenya and Uganda. PLOS Med 13(8):1–17. https://doi.org/10.1371/journal.pmed.1002099

Bauch C (2002) A versatile ODE approximation to a network model for the spread of sexually transmitted diseases. J Math Biol 45(5):375–395. https://doi.org/10.1007/s002850200153

Bersini H (2012) UML for ABM. J Artif Soc Soc Simul 15(1):Article 9. http://jasss.soc.surrey.ac.uk/15/1/9.html. Accessed 15 Sept 2018

Binmore K (2007) Playing for real: a text of game theory. Oxford University Press, Oxford

Blower SM, Boe C (1993) Sex acts, sex partners, and sex budgets: implications for risk factor analysis and estimation of HIV transmission probabilities. J Acquir Immune Defic Syndr 6(12):1347–1352

Caraël M, Glynn J, Lagarde E, Morison L (2004) Sexual networks and HIV in four African populations: the use of a standardised behvioural survey with biological markers. In: Morris M (ed) Network epidemiology: a handbook for survey design and data collection, chap 3. Oxford University Press, Oxford, pp 58–84

Chen MI, Ghani AC, Edmunds J (2008) Mind the gap: the role of time between sex with two consecutive partners on the transmission dynamics of gonorrhea. Sex Transm Dis 35(5):435–444

Delva W, Meng F, Beauclair R, Deprez N, Temmerman M, Welte A, Hens N (2013) Coital frequency and condom use in monogamous and concurrent sexual relationships in Cape Town, South Africa. J Int AIDS Soc 16:1–8

Deuchert E, Brody S (2007) Plausible and implausible parameters for mathematical modeling of nominal heterosexual HIV transmission. Ann Epidemiol 17(3):237–44

Dietz K (1988) On the transmission of HIV. Math Biosci 90:397–414

Eaton J (2013) The spread and control of HIV in southern Africa. Ph.D thesis

Eaton JW, Hallett TB, Garnett GP (2011) Concurrent sexual partnerships and primary hiv infection: a critical interaction. AIDS Behav 15(4):687–692. https://doi.org/10.1007/s10461-010-9787-8

Epstein H (2007) The invisible cure. Farrar, Strauss and Giroux, New York

Epstein H (2008) AIDS and the irrational. BMJ 337. http://www.bmj.com/content/337/bmj.a2638. Accessed 15 Sept 2018

Epstein H (2010) The mathematics of concurrent partnerships and HIV: a commentary on Lurie and Rosenthal 2009. AIDS Behav 14(1):29–30

Epstein H, Morris M (2011) Concurrent partnerships and HIV: an inconvenient truth. J Int AIDS Soc 14:1–11

Fu R, Gutfraind A, Brandeau ML (2016) Modeling a dynamic bi-layer contact network of injection drug users and the spread of blood-borne infections. Math Biosci 273:102–113. https://doi.org/10.1016/j.mbs.2016.01.003

Gaydosh L, Reniers G, Helleringer S (2013) Partnership concurrency and coital frequency. AIDS Behav 17(7):2376–2386

Goodreau S, Cassels S, Kasprzyk D, Montano D, Greek A, Morris M (2012) Concurrent partnerships, acute infection and HIV epidemic dynamics among young adults in Zimbabwe. AIDS Behav 16:312–322

Halperin DT, Epstein H (2004) Concurrent sexual partnerships help to explain Africa's high HIV prevalence: implications for prevention. Lancet 364(9428):4–6

Halperin DT, Epstein H (2007) Why is HIV prevalence so severe in southern Africa? South Afr J HIV Med 8(1):19–25

Harrison A, O'Sullivan L (2010) In the absence of marriage: long-term concurrent partnerships, pregnancy, and HIV risk dynamics among South African young adults. AIDS Behav 14:991–1000

Harrison A, Cleland J, Frohlich J (2008) Young people's sexual partnerships in KwaZulu-Natal, South Africa: patterns, contextual influences, and HIV risk. Stud Fam Plan 39(4):295–308
Hollingsworth TD, Anderson RM, Fraser C (2008) HIV-1 transmission, by stage of infection. J Infect Dis 198(5):687–693. https://doi.org/10.1086/590501
Hudson CP (1993) Concurrent partnerships could cause AIDS epidemics. Int J STD AIDS 4(5):249–253
Jenness SM, Biney AA, Ampofo WK, Nii-Amoo Dodoo F, Cassels S (2015) Minimal coital dilution in Accra, Ghana. J Acquir Immune Defic Syndr 69(1):85–91
Jones CI, Klenow PJ (2016) Beyond GDP? welfare across countries and time. Am Econ Rev 106(9):2426–2457. https://doi.org/10.1257/aer.20110236
Kharsany ABM, Frohlich JA, Yende-Zuma N, Mahlase G, Samsunder N, Dellar RC, Zuma-Mkhonza M, Abdool Karim SS, Abdool Karim Q (2015) Trends in HIV prevalence in pregnant women in rural South Africa. J Acquir Immune Defic Syndr 70(3):289–295. https://doi.org/10.1097/QAI.0000000000000761
Kim JH, Riolo RL, Koopman JS (2010) HIV transmission by stage of infection and pattern of sexual partnerships. Epidemiology 21(5):674–684
Kretzschmar M (1995) Deterministic and stochastic pair formation models for the spread of sexually transmitted diseases. J Biol Syst 03(03):789–801. https://doi.org/10.1142/S0218339095000721
Kretzschmar M, Dietz K (1998) The effect of pair formation and variable infectivity on the spread of an infection without recovery. Math Biosci 148(1):83–113. https://doi.org/10.1016/S0025-5564(97)10008-6
Kretzschmar M, Heijne JC (2017) Pair formation models for sexually transmitted infections: a primer. Infect Dis Model 2(3):368–378. https://doi.org/10.1016/j.idm.2017.07.002
Kretzschmar M, Morris M (1996) Measures of concurrency in networks and the spread of infection disease. Math Biosci 133:165–195
Kretzschmar M, White RG, Caraël M (2010) Concurrency is more complex than it seems. AIDS 24(2):313–315. https://doi.org/10.1097/QAD.0b013e328333eb9d
Lurie MN, Rosenthal S (2010a) The concurrency hypothesis in sub-Saharan Africa: convincing empirical evidence is still lacking. Response to Mah and Halperin, Epstein, and Morris. AIDS Behav 14(1):34–37
Lurie MN, Rosenthal S (2010b) Concurrent partnerships as a driver of the HIV epidemic in sub-Saharan Africa? the evidence is limited. AIDS Behav 14(1):17–24
Mah T, Halperin D (2010) Concurrent sexual partnerships and the HIV epidemics in Africa: evidence to move forward. AIDS Behav 14(1):11–16. https://doi.org/10.1007/s10461-008-9433-x
McCreesh N, O'Brien K, Nsubuga RN, Shafer LA, Bakker R, Seeley J, Hayes RJ, White RG (2012) Exploring the potential impact of a reduction in partnership concurrency on HIV incidence in rural Uganda: a modeling study. Sex Transm Dis 39(6):407–413. https://doi.org/10.1097/OLQ.0b013e318254c84a
Morris M, Kretzschmar M (1997) Concurrent partnerships and the spread of HIV. AIDS 11(5):641–648
Morris M, Kretzschmar M (2000) A microsimulation study of the effect of concurrent partnerships on the spread of HIV in Uganda. Math Popul Stud 8(2):109–133
Morris M, Epstein H, Wawer M (2010) Timing is everything: international variations in historical sexual partnership concurrency and HIV prevalence. PLoS One 5(11):1–8. https://doi.org/10.1371/journal.pone.0014092
Powers KA, Hoffman IF, Ghani AC, Hosseinipour MC, Pilcher CD, Price MA, Pettifor AE, Chilongozi DA, Martinson FEA, Cohen MS, Miller WC (2011) Sexual partnership patterns in Malawi: implications for HIV/STI transmission. Sex Transm Dis 38(7):657–666
Reniers G, Tfaily R (2008) Polygyny and HIV in Malawi. Demogr Res 19:1811–1830
Reniers G, Tfaily R (2010) An inquiry into the mechanisms linking polygyny, partnership concurrency and HIV transmission in sub-Saharan Africa. Technical report. Princeton University, Office of Population Research
Reniers G, Tfaily R (2012) Polygyny, partnership concurrency and HIV transmission in sub-Saharan Africa. Demography 49:1075–1101
Reniers G, Watkins S (2010) Polygyny and the spread of HIV in sub-Saharan Africa: a case of benign concurrency. AIDS 24(2):299–307
Sawers L (2013) Measuring and modelling concurrency. J Int AIDS Soc 16(1):17431. https://doi.org/10.7448/IAS.16.1.17431
Sawers L, Isaac AG (2017) Partnership duration, concurrency, and HIV in sub-saharan africa. Afr J AIDS Res 16(2):155–164. https://doi.org/10.2989/16085906.2017.1336105

Sawers L, Stillwaggon E (2010) Concurrent sexual partnerships do not explain the HIV epidemics in Africa: a systematic review of the evidence. J Int AIDS Soc 13(1):Article 34, 1–23. https://doi.org/10.1186/1758-2652-13-34

Sawers L, Isaac AG, Stillwaggon E (2011) HIV and concurrent sexual partnerships: modelling the role of coital dilution. J Int AIDS Soc 14(1):Article 44, 1–9. https://doi.org/10.1186/1758-2652-14-44

Tanser F, Bärnighausen T, Hund L, Garnett GP, McGrath N, Newell ML (2011) Effect of concurrent sexual partnerships on rate of new HIV infections in a high-prevalence, rural South African population: a cohort study. Lancet 378(9787):247–255

Watts CH, May RM (1992) The influence of concurrent partnerships on the dynamics of HIV/AIDS. Math Biosci 108(1):89–104

Welch G, Chick S, Koopman J (1998) Effect of concurrent partnerships and sex-act rate on gonorrhea prevalence. Simulation 71(4):242–249. https://doi.org/10.1177/003754979807100404

Young A (2005) The gift of the dying: the tragedy of aids and the welfare of future African generations. Q J Econ 120(2):423–466. https://doi.org/10.1093/qje/120.2.424