

Coevolution of risk perception, sexual behaviour, and HIV transmission in an agent-based model

Stephen Tully^{1,*}, Monica Cojocaru¹, Chris T. Bauch¹

July 4, 2014

1: Department of Mathematics and Statistics

University of Guelph

50 Stone Road East

Guelph, ON, N1G 2W1

Canada

Phone: +1 519-824-4120

Fax: +1 519-837-0221

Email addresses: stully@uoguelph.ca, mcojocar@uoguelph.ca, cbauch@uoguelph.ca

* Author for correspondence

for submission to *The Journal of Theoretical Biology*

Abstract

Risk perception shapes individual behaviour, and is in turn shaped by the consequences of that behaviour. Here we explore this dynamic in the context of human immunodeficiency virus (HIV) spread. We construct a simplified agent-based model based on a partner selection game, where individuals are paired with others in the population, and through a decision tree, agree on unprotected sex, protected sex, or no sex. An individual's choice is conditioned on their HIV status, their perceived population-level HIV prevalence, and the preferences expressed by the individual with whom they are paired. HIV is transmitted during unprotected sex with a certain probability. As expected, in model simulations, the perceived population-level HIV prevalence climbs along with actual HIV prevalence. During this time, HIV- negative individuals increasingly switch from unprotected sex to protected sex, HIV+ individuals continue practicing unprotected sex whenever possible, and unprotected sex between HIV+ and HIV- individuals eventually becomes rare. We also find that the perceived population-level HIV prevalence diverges according to HIV status: HIV- individuals develop a higher perceived HIV prevalence than HIV+ individuals, although this result is sensitive to how much information is derived from global versus local sources. This research illustrates a potential mechanism by which distinct groups, as defined by their sexual behaviour, HIV status, and risk perceptions, can emerge through coevolution of HIV transmission and risk perception dynamics.

Keywords: game theory, behaviour-prevalence model, HIV transmission model, agent-based model, epidemiology

1 Introduction

The World Health Organization (WHO) has classified HIV as a pandemic; the infection continues to spread in most nations, and is particularly prevalent in some of the world's poorest nations [1]. Medical treatment has reduced the annual rate of AIDS cases dramatically, however, the rate of new HIV infections has not dropped significantly [2]. Infection occurs by transfer of blood, semen, vaginal fluid, or breast milk [1], but in many countries the primary mode of transmission is sexual contact. For example, it is estimated that in the United States, 80% of new HIV cases are a result of unprotected sex [3]. Much of this infection is driven by so-called "core groups" that are responsible for a disproportionate amount of infection, and can drive the epidemic in some populations [4].

How individuals choose sexual partners and how they perceive their risk of infection within those partnerships appear to be major determinants of sexually risky behaviour as well as HIV infection risk [5,6]. For example, in men who have sex with men (MSM), a lower perceived risk of infection within a partnership is correlated with a higher probability of engaging in unsafe sex [6]. Another study in MSM found that, although most HIV-positive males practiced safer sex with HIV-negative or HIV-status-unknown partners, having more lifetime male sexual partners was correlated with being more likely to practice unsafe sex with such partners [7]. Hence, risk perception, safe sex practices, and sexual partner selection are determinants of HIV transmission between sex partners [8], and thus HIV prevalence in the population.

However, HIV prevalence in turn determines risk perception and sexual behaviour. For example, evidence from African HIV epidemics suggests that individuals adopt safer sex practices in response to rising HIV prevalence [9]. Individuals also appear to respond to how severe they expect HIV infection to be, with evidence that both HIV+ and HIV- individuals engage in more unprotected sex when anti-retroviral therapy (ART) is available; this occurs because ART causes HIV to be perceived to be more of a chronic disease than a fatal disease [10]. These findings highlight the feedback loop between risk perception and sexual partner selection/safe sex practices on the one hand, and HIV prevalence on the other hand: when HIV prevalence is low, individuals perceive less infection risk and engage in riskier

sexual behaviour, but that in turn can increase HIV prevalence. When HIV prevalence is high, individuals change their sexual behaviour to protect themselves.

In the field of behavioural epidemiology, models of behavioural influences on HIV transmission were some of the first to be developed. These approaches have analyzed, for example, the impact of contact patterns between various groups and other population heterogeneities, and how these factors interact with behavioural effects pertaining either to adherence to interventions or sexual behaviour, where behavioural change is either imposed through changes in a fixed parameter, or where it responds in a more fluid way to changes in HIV prevalence, as a model variable or a prevalence-dependent parameter [9, 11–20].

A few of these models have focussed specifically on the topic of sexual partner selection and how that influences HIV transmission [11, 18]. Game theory provides a very natural framework for addressing such topics [18], since it allows individual preferences to be concisely described and provides a way of determining the collective outcomes of these individual-level strategic decision-making processes [21] (as opposed to modelling the effects of behaviour in a more phenomenological way, such as a functional form describing how the person-to-person transmission rate declines with increasing prevalence). However, it also requires assuming that individuals will maximize their expected utility given the preferences and likely actions of others. In the context of partner selection and HIV transmission, the ‘game’ is a sexual interaction between two partners, and the payoff is the utility outcome of a player’s choice of unprotected or safe sex, given what they believe is the HIV status of their partner [18].

In this paper, we develop an agent-based simulation model that extends a previous game theoretical model of sexual partner selection and HIV transmission [18]. Individual perception of HIV infection risk evolves a function of the individual’s past sexual encounters, and individuals decide on either unprotected sex, protected sex, or no sex with the individuals with whom they are paired in each time step. Instead of exogenously imposing a pre-existing population structure where individuals engage in either high-risk or low-risk behaviour, our objective is to explore how distinct population groups—as defined by differing risk perception, HIV status, and sexual behaviours—can emerge from the coevolution of individual-level risk perception and HIV transmission, through partner selection games. We also wish to gain insight into potential mechanisms for how assortative sexual mixing with respect to level of

riskiness in sexual behaviour might emerge from this coevolutionary process.

2 Methods

The partner selection model is based on the dynamic signalling game described in Ref. [18], except modifications are made regarding how individuals perceive their risk of HIV infection as a function of their encounters, and how this perception is carried into future rounds of the game. A game represents a strategic interaction between two or more individuals. A dynamic signalling game is a type of sequential game with two players passing signals back and forth [18, 21, 22]. The players may be of differing types. Only the sender knows his/her own type: the receiver observes the signal but not the type of the sender. The receiver proceeds to choose an action while updating their beliefs regarding the type of the other player.

2.1 Partner selection model

The population consists of N individuals forming a highly active core group population of MSM, such as the patrons of an urban bathhouse. The population is not structured with respect to age, gender, or other forms of social/cultural heterogeneity. In each round, individuals are paired at random to play a “Risky Sex Game”. One of the pair is randomly assigned to be actor 1 (the initiator) and the other is assigned to be actor 2. The strategy options are Unprotected Sex (US), Protected Sex (PS) and No Sex (NS). Individual status is either HIV+ or HIV-, hence there are three possible pairs in each round: two HIV-, two HIV+ and HIV+/HIV-. The interaction describes a potential casual sexual encounter between two actors who may or may not have had an interaction before. Each actor knows whether or not they carry HIV, but do not know the HIV status of those they are paired with. Actors hold their own HIV status as private. Everyone knows the sexual preferences of HIV+ and HIV- actors, and based on their sexual preferences, an actor will alter their belief about the HIV status of the individual they are paired with. HIV- actors will not have unprotected sex with those that they believe are HIV+, while HIV+ actors prefer unprotected sex regardless of HIV status of partner. Both actors have a belief regarding the HIV status

of their partner, but they do not know with certainty and do not find out the truth, even if they eventually have sex (unless they become infected). This game will be embedded in an agent-based simulation model.

After being paired, the actors choose whether to play the Risky Sex Game with their partner with probability

$$\rho(c) = \frac{\omega + c}{1 + c} \quad (1)$$

where $\rho(c)$ represents the probability that they will play the game, c represents the cumulative sex acts that have occurred between actor 1 and actor 2 until this point in time, and the “gregariousness” parameter ω controls the how the probability of the game being played depends on how often the actors have engaged in sex previously. Equation (1) partially captures the effect of learning and reputation. In our baseline analysis, we assume $\omega = 1$, hence $\rho = 1$, but we explore $\omega < 1$ in sensitivity analysis. If the actors decide not to play the game this round, they must wait until the next round of the game to be re-paired.

Should the actors decide to interact further, the game will commence as follows:

1. Actor 1 offers to have either protected sex (PS) or unprotected sex (US).
2. Actor 2 either accepts or makes a counter offer of PS or US.
3. Actor 1 either accepts final offer or ends the interaction with no sex (NS).

The decisions available to each individual in each step of the Risky Sex Game are depicted in Supplementary Figure 1.

A utility is a value assigned to a possible outcome representing how an individual ranks their preferences. The utility U_i for actor i of outcome k ($k=US, PS, NS$) when the HIV status of actor i is s_i ($s_i=+, -$) and the HIV status of actor j with whom they are paired is s_j ($s_j=+, -$) is given by

$$U_i(k, s_i, s_j) \quad (2)$$

where the utilities for various outcomes as they depend on HIV status follow the inequalities contained in Table 1, which also contains the baseline utility values.

Because actors do not actually know the HIV status of their partner, each actor accrues an expected utility that depends on how strongly they believe their partner is HIV+. The

expected utility is a way of weighting utilities with an individual's inherent belief of HIV prevalence in the population, as represented by the b -value. The expected utility E_i for actor i in a given round when the outcome is k is therefore given by

$$E_i(b_t, s_i) = b_t U_i(k, s_i, +) + (1 - b_t) U_i(k, s_i, -) \quad (3)$$

where b_t is the probability with which actor i believes that their partner is HIV+, in the current round t of the game. With probability b_t , actor i believes that their partner is HIV+ and hence actor i would receive utility $U_i(k, s_i, +)$. Otherwise, actor i would receive utility $U_i(k, s_i, -)$. It is assumed that actors will maximize their personal expected utility.

The quantity b_t is specific to each actor (agent), and varies over the course of the simulation as actors update their b values based on their encounters. The value of b_t can be updated during a round, as well as between rounds. At the start of a round, actor 1 offers actor 2 their optimal choice of outcome k based on equation (3), using their current value of b_t based on previous interactions. If actor 1 offers US, then actor 2 increases his/her b -value according to

$$b_t \mapsto \alpha + (1 - \alpha)b_t \quad (4)$$

where α is called the ‘‘historical influence parameter’’, since it measures how much of b_t is from past encounters versus the current offer. If α is high, then actor 2's assessment of actor 1's HIV status is strongly dependent on the fact that actor 1 made an offer of unprotected sex, whereas if α is low, actor 1's offer had little impact on actor 2's evaluation of whether or not actor 1 is HIV+. On the other hand, if actor 1 offers PS, then actor 2 decreases his/her b -value according to

$$b_t \mapsto (1 - \alpha)b_t \quad (5)$$

Actor 2 then evaluates their new expected utility from equation (3), using the updated value of b_t , and responds to actor 1 with their own optimal outcome. If the two choices of outcome are in agreement, the actors will engage in that type of sex (US or PS). However, if they are in disagreement, Actor 1 will receive Actor 2's counter-proposal, and Actor 1 will update their b_t value in the same way, using equations (4) and (5) depending on the counter-offer. Actor 1 then uses this updated b_t value to determine their new expected utility, again according to equation (3). If the optimal outcome for Actor 1 has changed and the actors are now in

agreement, then they engage in that optimal choice of sex. If they remain in disagreement, then outcome is no sex (NS).

The value of b_t is also updated between rounds, based on information the actors are assumed to be able to receive regarding the actual prevalence of HIV in their population. This updating follows the equation

$$b_t = \beta b_{t-1} + (1 - \beta) \frac{I_t}{N} \quad (6)$$

where I_t is the prevalence of HIV in the population at the start of round t and β is the “local information” parameter governing the extent to which an actor’s b_t value is based on information from their interactions versus globally available, accurate information on HIV prevalence in their population. b_t is the value a player carries into round t , based on their b_t value at the end of the previous round $t - 1$ and how it has been modified through equation (6) inbetween rounds. For our baseline analysis we assume $\beta = 1$ and we explore $\beta < 1$ in sensitivity analysis. In general, we expect β to be close to 1 since the timescale of an HIV epidemic in a small core group population can be quite rapid [23,24].

At the end of each round, each pair breaks up and individuals are paired randomly again for the next round. Each player draws their initial b-value b_0 randomly from a normal distribution of mean $\mu = 0.05$ and standard deviation $\sigma = 0.1$; sampled values less than 0 or greater than 1 were discarded. We chose a relatively small value for the initial b-value because recruited individuals in the model are assumed to be younger, and both the perceived prevalence of HIV in a population as well as the perceived risk of contacting HIV or other STI’s are less in younger individuals [25–27]. Table 2 describes the different parameters used in the model and also provides their baseline values.

2.2 Transmission and natural history model

We assume a time scale on the order of a single HIV outbreak in a core group population. Hence, we neglect recruitment and death, and assume a Susceptible-Infected (SI) natural history where individuals are initially susceptible, and remain infectious indefinitely once they are infected. Each round of the game corresponds to a timestep of the simulation model, and each timestep corresponds to approximately three days. If an outcome of Unprotected

Sex occurs between an HIV+ and an HIV- actor, then HIV is transmitted with probability $\tau = 0.01$. For simplicity we assume Protected Sex is perfectly effective in preventing transmission [28]. The population size is $N = 1000$ players, with 5% of the population being infected at $t = 0$.

3 Results

3.1 Baseline scenario

The time series of HIV prevalence ($I(t)$) initially shows a typical Susceptible-Infected natural history outbreak pattern, with prevalence increasing exponentially quickly in the early stages. However, instead of infecting the vast majority of the population, HIV prevalence subsequently approaches an equilibrium \bar{I} comprising approximately 60% of the population (Figure 1a).

The approach to \bar{I} is asymptotic, and this stabilization occurs because HIV+ and HIV- actors evolve differing b-values in the course of the outbreak. For the first 50 time steps, the average b value of HIV+ actors, $\langle b+ \rangle$, increases rapidly and at roughly the same rate as the average b value of HIV- actors, $\langle b- \rangle$. However, divergence between $\langle b+ \rangle$ and $\langle b- \rangle$ emerges suddenly after 50 time steps: $\langle b- \rangle$ continues to climb with increasing HIV prevalence while $\langle b+ \rangle$ stabilizes. Before long, $\langle b+ \rangle$ and $\langle b- \rangle$ have diverged considerably (Figure 1a). The heightened b value of HIV- actors means they often assume those who offer them US are HIV+ and hence they respond with a counteroffer of PS, avoiding HIV infection. Hence, HIV- actors are protected from infection by their high b-values, and this slows down the transmission of HIV, resulting in stabilization of HIV prevalence at approximately 60%. By time step 1500, the population is divided into two groups: one group is HIV+ and tends to have a low perception of HIV prevalence in the population, and the other groups is HIV- and tends to have a higher perception of HIV prevalence in the population.

The divergence of $\langle b- \rangle$ and $\langle b+ \rangle$ occurs primarily because actor b-values are updated only when a counter-offer is called for. As HIV prevalence increases, an HIV- Actor 2 will increase his b-value when paired with an HIV+ Actor 1 offering US, but not when paired with an HIV- Actor 1 offering PS. Hence, as HIV prevalence increases, HIV- actors are increasingly

being offered US, and they conclude that HIV prevalence is increasing. Similarly, an HIV+ Actor 2 will decrease his b-value when paired with an HIV- Actor 1 offering PS, but not when paired with an HIV+ Actor 1. Hence, the b-value for HIV+ actors is lower than for HIV- actors. A modification of the model allowing actors to update their b-values even when the preferences of the actors are in agreement shows that $I(t)$ and $\langle b \rangle$ are similar, but there is less divergence in $\langle b- \rangle$ and $\langle b+ \rangle$ (Supplementary Figure 2).

The divergence is also driven by a positive feedback loop: actors with slightly higher b-values in the beginning are somewhat more likely to avoid HIV infection, which in turn drives their b-value to higher levels according to the mechanism described in the previous paragraph. Likewise, actors with slightly lower b-values are more likely to become infected, and their lower b-value will be reinforced by the same mechanism.

A stratification of outcomes by HIV status of actors provides more detail on how these two groups emerge and interact (Table 3). In the early stages of the outbreak when $I(t)$ is still small, the most common type of outcome for all three types of pairs is Unprotected Sex. HIV+/+ pairs always engage in Unprotected Sex, and HIV-/- or HIV+/- pairs engage in Unprotected Sex significantly more often than Protected Sex or No Sex. However, by the time that HIV is highly prevalent, the most common outcome is Unprotected Sex in HIV+/+ pairs. Unprotected Sex hardly ever happens in HIV+/- or HIV-/- pairs, where Protected Sex is instead the most common outcome. In summary, when HIV is highly prevalent, HIV+ actors engage in Unprotected Sex with other HIV+ actors but Protected Sex with HIV- actors, whereas HIV- actors engage in Protected Sex with both HIV+ and HIV- actors, or oftentimes No Sex if the other actor is HIV+. Hence, by this stage, the population has assorted into two distinct groups with very distinct perception and sexual practices.

As expected, HIV prevalence at the pseudo-equilibrium \bar{I} and the average b-value $\langle b \rangle$ are insensitive to changes in the utility for protected sex, u_{PS} (Figure 2), since changes in the range $0 < u_{PS} < 100$ do not impact the relative utility rankings (Table 1). Increasing the transmission probability τ increases both the HIV prevalence at the pseudo-equilibrium as well as $\langle b \rangle$ (Figure 2). Changes in the historical influence parameter α also do not have much impact on \bar{I} and $\langle b \rangle$ unless $\alpha = 0$, in which case actors' b-values are fixed and actors never use information from their interactions to modify their b-value (Figure 2). The differentiation in

b-values between HIV+ and HIV- individuals is also very robust to changes in u_{PS} , τ , and α (Figure 3).

3.2 Scenario $\beta < 1$

Our baseline assumption is $\beta = 1$, meaning that individuals do not have access to a global source of accurate information on HIV prevalence. When $\beta < 1$, individuals base their b-value partly on global HIV prevalence (Equation (6)). Average HIV prevalence at the pseudo-equilibrium \bar{I} and the average b-value $\langle b \rangle$ are relatively unaffected by changes in β , however, the divergence between $\langle b- \rangle$ and $\langle b+ \rangle$ disappears as β decreases, since individuals increasingly base their perception of HIV prevalence on actual prevalence in the population, rather than on their interactions with others (Supplementary Figures 3,4). This is robust to changes in α , u_{PS} and τ (Supplementary Figures 5, 6, Supplementary Table 1). We speculate that higher values of β close to 1 are most realistic, since low values of β correspond to having accurate information on global HIV prevalence on a weekly basis, which is not realistic for most populations. We also note that the rise in $\langle b \rangle$, $\langle b- \rangle$, and $\langle b+ \rangle$ is more gradual than the baseline scenario, indicating that Unprotected Sex will be preferred by both types for a longer amount of time.

3.3 Scenario $\omega < 1$

Our baseline assumption is $\omega = 1$, meaning that actors play the Risky Sex Game regardless of the previous sexual history of those they are paired with (Equation (1)). The divergence between $\langle b- \rangle$ and $\langle b+ \rangle$ persists when $\omega < 1$, and \bar{I} and $\langle b \rangle$ are likewise relatively unaffected (Supplementary Figures 7-9). The population still evolves into two groups characterized by divergent HIV status and sexual behaviour (Supplementary Table 2). The main impact of decreasing ω is to cause system evolution to slow down on account of a larger number of rejected pairings due to lack of a previous history between actors. However, long-term dynamics are unchanged. We expect larger values of c to apply in populations that are smaller, where individuals are more likely to be known to one another.

3.4 Sensitivity analysis

We analyzed a variant of the model where individuals do not immediately know whether they have been infected, and only discover it if they are tested. Testing occurs with probability t per timestep and we assume 100% accuracy. HIV+ individuals only change their utility from HIV- to HIV+ once they become aware of their changed HIV status. When the testing probability t is sufficiently high, then the predictions of this variant model agree with those of the baseline model (Supplementary Figures 4, 5). As t decreases, some differences begin to emerge: total HIV prevalence is higher when the testing probability is very low (Supplementary Figure 4). This occurs because there are more individuals in the population who are HIV+ and yet make decisions according to HIV- utilities, unwittingly spreading infection in cases where the outcome of the interaction is unprotected sex with an HIV- individual. Also, as t decreases, the separation in b-values between HIV+ and HIV- individuals persists even as t approaches zero (Supplementary Figure 5). However, when t is high, the b-value of HIV+ individuals who are not aware of their status aligns with the b-value of HIV+ individuals who are aware of their status, whereas when t is low, the b-value of HIV+ individuals who are not aware of their status aligns with the b-value of HIV- individuals (Supplementary Figure 5).

We assumed a population size of 1,000 because the number of clientele of a typical urban bathhouse is on this order of magnitude [29,30]. However, we also experimented with varying N from 500 to 10,000, finding that model predictions are close to our baseline population of $N = 1,000$ (the $N = 5,000$ case is shown in Supplementary Figure 6).

We also conducted a probabilistic sensitivity analysis, where distributions are defined around certain input parameters based on lowest and highest plausible values, repeated model simulations are conducted after sampling each input parameter from the distribution independently for each simulation, and mean and standard deviation of the resulting outputs are reported. This was conducted on $I_{initial}$, τ , α , and b_0 . Results can be seen in tables 4 and 5, and indicate a continued separation between average b-values for HIV+ and HIV - individuals. These data demonstrates the robustness of the model with respect to uncertainty in those four parameters.

4 Discussion

Here we have explored a potential mechanism for how distinct groups—as defined by risk perception, HIV status, and sexual behaviour—can emerge out of the coevolution of HIV transmission and risk perception through sexual partner dynamics. The changing patterns of risk perception and sexual behaviour are plausible: the perceived level of HIV prevalence in the population rises along with the actual HIV prevalence, and as HIV spreads, more HIV- individuals opt for protected sex.

We observed that the actors’ b-values (the perceived probability that a randomly chosen individual in the population is HIV positive) diverged according to HIV status, with HIV- actors developing higher b-values and HIV+ actors developing lower b-values. This divergence was robust to changes in all model parameters except for β , which controlled the extent to which information on HIV prevalence is taken from global sources versus being based on their interactions with others. However, divergence was found to be less pronounced if actors were allowed to update their b-values not only when a counter-offer is discordant but also when the other actor agrees to what has been offered.

In at least one population, patterns of risk perception and sexual behaviour similar to those predicted by our model have been found. In a study of MSM in the United States, it was found that study participants with unrecognized HIV infection perceived themselves to have a low lifetime risk for acquiring HIV (i.e., they had low b-values), and also engaged in riskier sex [26]. This is consistent with our model prediction that HIV+ individuals have a lower perceived risk of being infected than HIV- individuals.

Because our objective was to identify a potential mechanism for how coevolution of risk perception and HIV transmission could lead to the emergence of groups, and because our focus was on a small, relatively homogeneous core group population over a relatively short timescale, we did not include important heterogeneities in transmission, population structure, or risk perception, although some of these could influence our results quantitatively or qualitatively. For example, we neglected age structure, gender structure, vital dynamics (birth and death) and the details of HIV natural history in infected hosts. Future work could explore whole-population models and also build greater realism into the decision-

making process. Our assumption that each actor knows whether or not they carry HIV is an idealization; in fact, there may be a significant lag between contracting HIV and becoming aware of HIV status, even in industrialized countries [31]. However, we explored the impact of this assumption in sensitivity analysis.

The model predicted that, at the pseudo-equilibrium, unprotected sex will usually be practiced in HIV+/HIV+ pairs, and protected sex in HIV+/HIV- or HIV-/HIV- pairs. This mixing pattern is the outcome of our assumed mechanisms of interaction between risk perception, partner selection, and infection transmission. Similar patterns of assortative mixing are observed in real populations. We do not claim that such assortative mixing patterns are sufficiently explained by this mechanism, however, we hypothesize that this mechanism can contribute to these mixing patterns.

HIV prevalence is not the only determinant of an individual's perception of HIV infection risks, and many studies consider correlates for perceived risk such as gender, socioeconomic status and age, which are also influential. These previous studies are generally cross-sectional in nature, whereas longitudinal studies would be better suited to quantify factors of a dynamic nature such as feedback from evolving HIV prevalence during an outbreak. We suggest efforts should be made to conduct longitudinal population surveys to better gauge how risk perception and sexual partner practices evolve in response to changes in HIV prevalence and through the partner selection process, since our model indicates that it plays a central role in whether the b-values of HIV+ and HIV- individuals diverge or not.

Previous modelling approaches have not always considered how risk perception and transmission dynamics coevolve in the case of HIV. Our model highlights how individual perception of personal HIV infection risk can coevolve with HIV population prevalence in the course of an epidemic. This implies that risk perception is not a fixed, unchanging quantity that can be adequately described out of the context of HIV prevalence in the study population. Rather, risk perception influences HIV prevalence, and HIV prevalence in turn influences risk perception, so that they form a single, coupled system. It is difficult for cross-sectional studies to capture this effect, which plays out over time. However, future longitudinal studies could test this by attempting to correlate perceived HIV infection risk with actual HIV prevalence at different time points in an HIV epidemic, as well as monitor how perception

of personal risks and local or global HIV prevalence changes before and after infection with HIV. This will become increasingly important as HIV vaccines become a reality over the next decade, since HIV vaccines might change how HIV infection risk and/or severity are perceived, and thus might change behaviour as well. We suggest that more attention needs to be given to this potential interaction, and more generally to the issue of the dynamics of core group formation and maintenance in the context of sexually-transmitted infections.

5 Acknowledgements

This work was supported by a CIHR Operating Grant and an NSERC Discovery Grant to C.T. Bauch, and an NSERC Discovery Grant to Monica Cojocaru.

References

- [1] World Health Organization. Health topics: HIV/AIDS, Accessed March 6th 2013. <http://www.who.int/bulletin/volumes/85/11/06-033779/en/index.html>.
- [2] Joint United Nations Programme on HIV/AIDS. UNAIDS 2011 world AIDS day report, 2011.
- [3] G. Marks, N. Crepaz, and S. Robert. Estimating sexual transmission of HIV from persons aware and unaware that they are infected with the virus in the USA. *AIDS*, 20(10):1447–1450, 2006.
- [4] F. Plummer, N.J. Nagelkerke, S. Moses, J. Ndinyaa-Achola, J. Bwayo, and E. Ngugi. The importance of core groups in the epidemiology and control of HIV-1 infection. *AIDS*, 5(1):69–76, 2004.
- [5] S. M. Blower, S. K. Service, and D. Osmond. Calculating the odds of HIV infection due to sexual partner selection. *AIDS and Behavior*, 1(4):273–274, 1997.
- [6] R. Remien, G. Wagner, C. Dolezal, and A. Carballo-Dieiguez. Factors associated with HIV sexual risk behavior in male couples of mixed HIV status. *Journal of Psychology and Human Sexuality*, 13(2):31–48, 2001.

- [7] N. Wenger, F. Kusseling, K. Beck, and M. F. Shapiro. Sexual behavior of individuals infected with the human immunodeficiency virus: The need for intervention. *Archives of internal medicine*, 154(16):1849–1854, 1994.
- [8] T. Barnett and P. Blaikie. *AIDS in Africa. Its present and future impact*. Guilford Press, New York, 1992.
- [9] T. Hallett, S. Gregson, O. Mugurungi, E. Gonese, and G. Garnett. Assessing evidence for behaviour change affecting the course of HIV epidemics: A new mathematical modelling approach and application to data from zimbabwe. *Epidemics*, 1(1):108–117, 2009.
- [10] N. H. T. M. Dukers, J. Goudsmit, J. B. F. de Wit, M. Prins, G. Weverling, and R. Coutinho. Sexual risk behaviour relates to the virological and immunological improvements during highly active antiretroviral therapy in HIV-1 infection. *AIDS*, 15(3):369–378, 2001.
- [11] J. Koopman, C. Simon, J. Jacquez, J. Jill, L. Sattenspiel, and T. Park. Sexual partner selectiveness effects on homosexual HIV transmission dynamics. *JAIDS*, 1(5):486–504, 1988.
- [12] S. Blower and A. McLean. Prophylactic vaccines, risk behavior change, and the probability of eradicating HIV in san francisco. *Science*, 265:1451–1454, 1994.
- [13] K. Hadeler and C. Castillo-Chavez. A core group model for disease transmission. *Mathematical Biosciences*, 128:41–55, 1995.
- [14] M. Kremer. Integrating behavioral choice into epidemiological models of AIDS. *Quarterly Journal of Economics*, 111(2):549–573, 1996.
- [15] J. Hyman and J. Li. Behavior changes in SIS STD models with selective mixing. *SIAM Journal of Applied Mathematics*, 57(4):1082–1094, 1997.
- [16] Y. Hsieh and S. Sheu. The effect of density-dependent treatment and behavior change on the dynamics of HIV transmission. *Journal of Mathematical Biology*, 43(1):69–80, 2001.

- [17] D. Greenhalgh and F. Lewis. Stochastic models for the spread of HIV amongst intravenous drug users. *Stochastic Models*, 17(4):491-512, 2001.
- [18] K. Schroeder and F. Rojas. A game theoretical analysis of sexually transmitted disease epidemics. *Rationality and Society*, 14(3):353-383, 2002.
- [19] L. Sattenspiel and C. Castillo-Chavez. Environmental context, social interactions, and the spread of HIV. *American Journal of Human Biology*, 2(4):397-417, 1990.
- [20] R. Smith and S. Blower. Could disease-modifying HIV vaccines cause population-level perversity? *The Lancet Infectious Diseases*, 4(10):636-639, 2004.
- [21] L. Thomas. *Games, Theory and Applications*. Dover Publications, New York, 2003.
- [22] T. Basar and G. Olsder. *Dynamic noncooperative game theory*. SIAM, 1999.
- [23] B. Vuylsteke and S. Jana. Reducing HIV risk in sex workers, their clients and partners. In Peter Lampthey and Helene D. Gayle, editors, *HIV/AIDS prevention and care in resource-constrained settings.*, pages 187-210. Family Health Institute, 2001.
- [24] F. Behets, J. Andriamiadana, D. Rasamilalao, N. Ratsimbazafy, D. Randrianasolo, G. Dallabetta, and M. Cohen. Sexually transmitted infections and associated socio-demographic and behavioural factors in women seeking primary care suggest Madagascar's vulnerability to rapid HIV spread. *Tropical Medicine and International Health*, 6(3):202-211, 2001.
- [25] A. Adedimeji. Beyond knowledge and behavior change: The social-structural context of HIV/AIDS risk perceptions and protective behavior among young urban slum inhabitants in Nigeria. *Takemi Program in International Health 2004-2005*, 2005.
- [26] D. MacKellar, L. Valleroy, G. Secura, S. Behel, T. Bingham, D. Celentano, B. Koblin, M. LaLota, D. Shehan, H. Thiede, and L. Torian. Perceptions of lifetime risk and actual risk for acquiring HIV among young men who have sex with men. *AIDS Behaviour*, 11(2):263-270, 2007.

- [27] M. Gardner and L. Steinberg. Peer influence on risk taking, risk preference, and risky decision making in adolescence and adulthood: an experimental study. *Developmental Psychology*, 41(4):625–635, 2005.
- [28] R. Gray, M.J. Wawer, R. Brookmeyer, N. Sewankambo, D. Serwadda, F. Wabwire-Mangen, T. Lutalo, X. Li, T. vanCott, T. Quinn, and The Rakai Project Team. Probability of HIV-1 transmission per coital act in monogamous, heterosexual, HIV-1-discordant couples in rakai, uganda. *The Lancet*, 357(9263):11491153, 2001.
- [29] G. Richwald, D. Morisky, G. Kyle, A. Kristal, M. Gerber, and J. Friedland. Sexual activities in bathhouses in Los Angeles county: implications for AIDS prevention education. *The Journal of Sex Research*, 25(2):169–180, 1988.
- [30] D. Binson, L. Pollack, J. Blair, and W. Woods. HIV transmission risk at a gay bathhouse. *The Journal of Sex Research*, 47(6):580–588, 2010.
- [31] R. Bayer and A. Fairchild. Changing the paradigm for HIV testing the end of exceptionalism. *The New England Journal of Medicine*, 355:647–649, 2006.

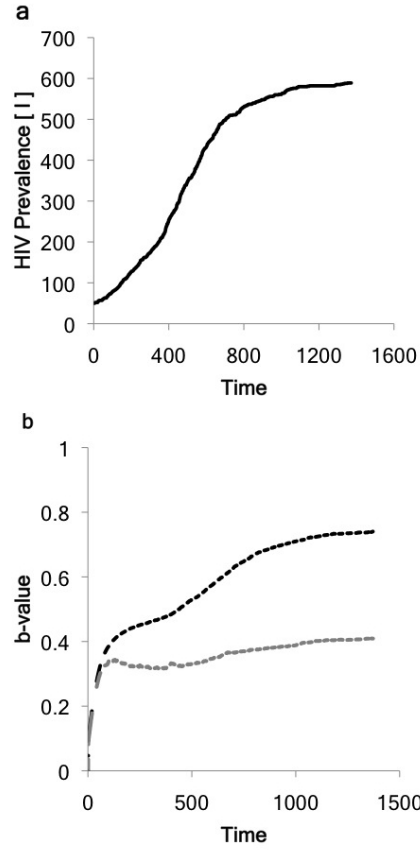


Figure 1: (a) HIV prevalence over time ($\langle I \rangle$, black). (b) Average b-values over time by type, HIV- (black dashed $\langle b- \rangle$), HIV+ (grey dashed $\langle b+ \rangle$). This illustrates the coevolution between the HIV prevalence in the population, and the assessment of risk among its individuals. We also note the divergence of b-values amongst the HIV- and HIV+ groups, demonstrating a significant difference in risk assessment.




Figure 2: HIV prevalence (% infected, $\langle I \rangle$, black), and total average b ($\langle b \rangle$, grey) for: (a) alpha, (b) transmission, (c) PS Utility. This demonstrates the robustness of the parameters used within the model. As expected, HIV prevalence rates are closely tied to average b -values amongst the population.




Figure 3: Average b-value for HIV- ($\langle b^- \rangle$, dashed black) and HIV+ ($\langle b^+ \rangle$, dashed grey) individuals for: (a) alpha, (b) transmission, (c) PS Utility. This also shows the robustness of the parameters by indicating the cluster separation remains for varying parameter values.

| s_i | s_j | x_{US} | x_{PS} | x_{NS} | preferences of actor i |
|------------------|------------------|----------------|----------------|----------------|--------------------------|
| actor i status | actor j status | utility for US | utility for PS | utility for NS | |
| HIV+ | HIV+ | 100 | 50 | 0 | $US > PS > NS$ |
| HIV+ | HIV- | 100 | 50 | 0 | $US > PS > NS$ |
| HIV- | HIV+ | -50 | 50 | 0 | $PS > NS > US$ |
| HIV- | HIV- | 100 | 50 | 0 | $US > PS > NS$ |

Table 1: Utilities for actor i , given status of actors i and j , and type of sex.

| Term | Definition | Baseline value |
|---------------|--|----------------|
| b_0 | Initial b-value: Initial probability value individually chosen from a beta distribution of mean 0.05 and a standard deviation of 0.1 | 0.05 |
| α | Historical influence parameter: Measures how much of the b-value will be based on past experience | 0.015 |
| τ | Transmission probability: Probability of HIV spread from an HIV+ player to an HIV- player through US | 0.01 [28] |
| u_{PS} | Protected Sex utility | 50 |
| β | Local information parameter: Proportion of the b-value that is a combination of an individual's interpretation of risk, and an accurate value of the actual proportion of infected individuals in the population | 1 |
| ω | Gregariousness parameter: Degree of willingness to interact with strangers | 1 |
| N | Population size | 1000 |
| $I_{initial}$ | Initial number of infected players at the beginning | 5% |
| I | Number of infected individuals in the population | 5% |

Table 2: Parameter definitions, and parameter values for baseline scenario.

| Time | | +/+ | +/- | -/- |
|------------|----|-------|-------|-------|
| 100 | US | 0.42 | 10.32 | 58.51 |
| Prev 8.2% | PS | 0 | 2.34 | 10.71 |
| | NS | 0 | 0.002 | 17.67 |
| 200 | US | 0.92 | 11.64 | 24.71 |
| Prev 11.7% | PS | 0 | 6.14 | 33.53 |
| | NS | 0 | 0.05 | 22.98 |
| 300 | US | 2.15 | 14.68 | 18.08 |
| Prev 16.9% | PS | 0 | 9.77 | 35.84 |
| | NS | 0 | 0.15 | 19.30 |
| 400 | US | 3.64 | 17.61 | 13.53 |
| Prev 22.6% | PS | 0 | 13.25 | 35.7 |
| | NS | 0 | 0.50 | 15.74 |
| 500 | US | 7.21 | 20.57 | 8.69 |
| Prev 31.5% | PS | 0 | 17.12 | 33.37 |
| | NS | 0 | 1.01 | 11.99 |
| 1500 | US | 34.96 | 1.44 | 0 |
| Prev 59.5% | PS | 0 | 25.64 | 16.52 |
| | NS | 0 | 21.34 | 0.07 |

Table 3: Cumulative sex acts for type. First three rows show total number of sex acts between $t = 0$ and $t = 100$, second three rows show total number of sex acts between $t = 100$ and $t = 200$, etc. $t = 1500$ represents cumulative sex acts between $t = 1400$ and $t = 1500$ at equilibria. This table shows evidence of how the preferred type of sex changes over time. As HIV prevalence and average b-value increases, we see a reduction in Unprotected Sex among HIV- couples as expected.

| Parameter | Distribution | Mean, SD |
|---------------|---------------------|---------------|
| α | Normal distribution | 0.015, 0.0025 |
| τ | Normal distribution | 0.01, 0.0025 |
| $I_{initial}$ | Normal distribution | 0.05, 0.015 |
| b_0 | Normal distribution | 0.05, 0.015 |

Table 4: Parameters adjusted in sensitivity analysis. Values recorded from normal distribution, with probabilities less than 0 or greater than 1 discarded.

| | Average | 2SD |
|-------------------------------|---------|--------|
| HIV Prevalence % | 0.5897 | 0.1081 |
| Average b-value | 0.5489 | 0.0439 |
| HIV - average b-value | 0.7403 | 0.0882 |
| HIV + average b-value | 0.4171 | 0.0609 |
| Average separation in b-value | 0.3232 | 0.0343 |

Table 5: Parameters adjusted in sensitivity analysis. Results continue to demonstrate significant divergence between average b-values between HIV + and HIV population.