On the Dynamics of Beliefs and Risky Sexual Behavior*

Flavio Toxvaerd†

January 13, 2017

Abstract. This paper analyzes the dynamic interplay between individuals' beliefs over transmission risks and their decisions on whether to engage in unprotected sexual behavior. The infectious disease is (possibly) asymptomatic, which makes the inference problem non-trivial. In general, the dynamics of beliefs can be non-monotone in time, even without shifts in actions. Both monogamous and concurrent relationships are considered. It is shown that if unprotected sex is voluntary, then a utilitarian social planner would prefer more unprotected sex than the level chosen in equilibrium. Furthermore, it is shown that altruism towards the partner will increase the amount of voluntary unprotected sex and thus increase the amount of infection.

Keywords: Asymptomatic infection, beliefs, infection dynamics, risky sexual behavior, learning, HIV/AIDS.

JEL Classification: D83, I12.

1. Introduction

The emergence and global spread of HIV/AIDS since the early nineteen eighties, is often referred to as the *invisible epidemic*. The reason is that infection with HIV can be asymptomatic for a very long time, allowing the disease to spread unnoticed in the population.[1] Thus, a central issue in understanding the rapid propagation and persistence of an easily avoidable and potentially fatal disease, is the fact that individuals may be infected for protracted periods of time, without being aware that they are infected. The median time between infection and the onset of symptoms (i.e. the incubation period) has been estimated to around ten years (see Bacchetti and Moss, 1989 and Chevret et al., 1992). The fact that HIV can remain asymptomatic for so long, seriously complicates the estimation of incidence and prevalence rates, as documented by Oster (2011).

Individuals appear to be fairly accurate in assessing their own probability of being infected (Delavande and Kohler, 2009 and 2011 and Dodds et al., 2004), although some evidence suggests that people tend to overestimate their probability of being infected (see Anglewicz and Kohler, 2009). At the population level, it is estimated that 18.1% of US HIV cases are undiagnosed. This number masks considerable variation across age groups, genders and ethnicity, ranging

*I wish to thank Larry Samuelson and David Levine for very useful conversations on the present work. I am also grateful for feedback from Francis Bloch, Bill Byrne, Diego Cerdeiro, Frederick Chen, Martin Cripps, Matt Elliott, Jeff Ely, Aytek Erdil, Alex Gershkov, Chryssi Giannitsarou, Ben Golub, Jeremy Greenwood, Chris Harris, Ed Kaplan, Shachar Kariv, Godfrey Keller, Konstantin Matthies, Stephen Morris, Alessandro Pavan, Theodore Papageorgiou, Nicola Pavoni, Brian Rogers, Richard Weber, Jossy Sayir, Jose Sheinkman and Xavier Vives. I also thank seminar participants at the University of York, the University of Surrey, the University of Cambridge, Birkbeck College at the University of London, the 2015 Conference for Research on Economic Theory and Econometrics, Chania, Greece, the 2015 Cambridge Infectious Diseases Meeting of Minds and the 2016 UECE Lisbon Meetings in Game Theory and Applications. I also thank the Paris School of Economics for generously hosting me while working on parts of the manuscript.

†Faculty of Economics, University of Cambridge. Email: fmot2@cam.ac.uk.

[1]Formally, *symptoms* are subjective, while *signs* are objective evidence of infection. In the present analysis, individuals are assumed to have no advantage over others in terms of assessing their actual infection status and hence an individual will have symptoms of infection if and only if it has signs of infection. The two terms are therefore interchangeable and I will therefore use the former throughout.

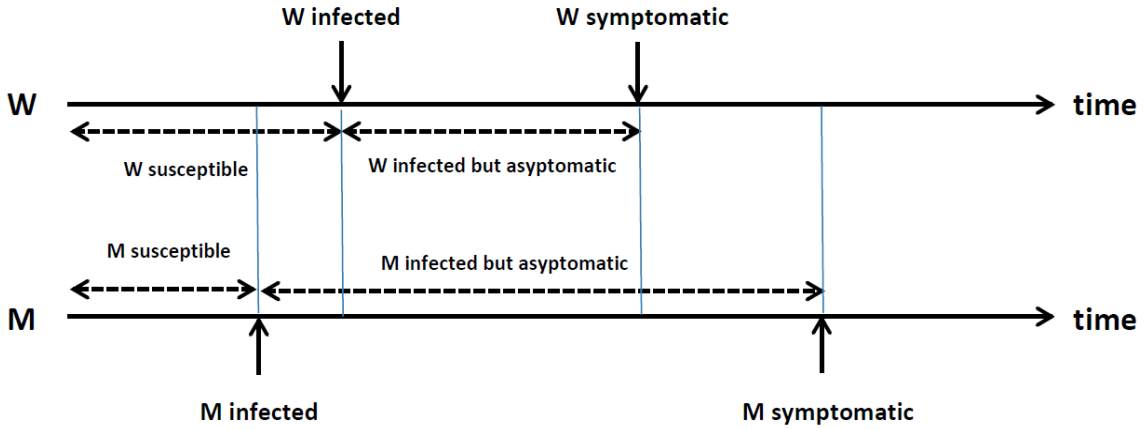

Figure 1: Timeline for the Man's and the Woman's Time of Infection and Onset of Symptoms.

from 10.6% for 55-64 year olds, to a whopping 59.5% for 13-24 year olds (see CDC, 2010). Data for the UK indicates that 24% of HIV cases are undiagnosed (see HPA, 2012).

In this paper, a simple model is presented in which individuals learn about their own and their partner's infection state by observing noisy information about these (i.e. they may observe symptoms if either or both individuals are infected). Furthermore, they decide whether to engage in protected or in unprotected sex. Learning and equilibrium decisions are shown to interact in interesting ways.

1.1. Symptoms and Inference. To appreciate the complexities involved in asymptomatic infection, consider the situation depicted in Figure 1. Assume that the latency period is negligible (i.e. that the time between infection and the onset of the disease is very short) and assume that the disease has a random incubation period (i.e. that the time between infection and the onset of symptoms is positive but random). The upper time line in Figure 1 shows the woman's date of infection and the subsequent onset date (i.e. the date at which symptoms appear). The lower time line shows the infection and onset dates for the man. By the nature of the problem, individuals would like to base their decisions on whether to engage in unprotected sex on the possible infection dates, but they cannot do so as they are not observed. At most, the individuals observe the onset dates and must use these (if available) to make inferences on which to base their decisions. From the perspective of the man in the example, the absence of symptoms in the woman does not imply the absence of infection, since he cannot distinguish between a susceptible and an infected but asymptomatic partner. Similarly, in this example, the woman shows symptoms before the man does. This raises the question as to whether the woman has been infected (exogenously) at a prior point in time, or she has become infected by the infected but still asymptomatic man (as is the case in the figure). One can therefore not conclude from the fact that the woman is the first to display symptoms, that she was the source of infection (or indeed that the man is infected at all). As each individual in the couple makes decisions on protected versus unprotected sex on the basis of the transmission risks, a key component of the individual's decision problem is to form beliefs about these.

1.2. Beliefs and Decisions. It is well documented that beliefs about the infection state of individuals and their partners interact in important ways with sexual behavior (see e.g. Akwara, 2003, Anglewicz and Kohler, 2009 and De Paula et al. 2013). The patterns of both sexual encounters and the use of condoms and other means to reduce the risk of disease transmission, are incredibly varied across populations, reflecting both differences in cultures, customs and faiths and differences in the costs of disease prevention and in the disease burden.

While it is thus difficult to generalize, it seems clear that substantial (double digit) proportions of individuals in long-term relationships either use condoms regularly, never use condoms or use condoms irregularly (see Williamson et al., 2006, Ellard, 2007 and Delva et al., 2013). Research has also shown that couples switch between protected and unprotected sex over the course of a relationship (see Fortenberry et al., 2002). To make progress in understanding such patterns of sexual decision making and to determine their interactions with aggregate disease dynamics, belief formation must thus be properly understood.

Existing contributions to the literature have taken a variety of approaches to beliefs. Philipson and Posner (1993) use a mean field approximation to both sides of a sexual encounter. This means in effect that not only does an individual believe its partner to be representative of the population, but it believes itself to be likewise representative (and implicitly, to ignore its own personal history of sexual decision making). With this assumption, the beliefs of the two individuals about themselves and each other are determined by aggregate characteristics. Greenwood et al. (2013) allow for updating about one's own health, but use mean field approximations for the different classes of potential partners (namely those seeking short-term unprotected sex, short-term protected sex and long-term unprotected sex). In their analysis, short-term sex (whether unprotected or protected) is always random, anonymous and one-off. While they do allow for long-term relationships, within such, individuals make no decisions and hence updating is incidental.[2] Hsu and Hsieh (2007) also use mean field approximations for all individuals in a classical compartmental model augmented with asymptomatically infected individuals, but they disregard decision making altogether and is thus not comparable to the present work. Similarly, Fisman et al. (2013) study the role of asymptomatic infection, but disregard both decision making and beliefs.

It should be emphasized that most (but not all) dynamic models of decision making under the threat of infection require some belief formation. This is particularly so with models of vaccination and treatment (see e.g. Geoffard and Philipson, 1997, Auld, 2003, Toxvaerd, 2012 and references therein). The important difference between that strand of literature and the work presented here, is that individuals form beliefs about aggregates, rather than about specific potential partners.[3] The assumptions maintained in those papers amount to very special implicit assumptions about interaction patterns in the population, which will be discussed further below.

Last, there is a connection between this paper and the testing literature, but I will postpone a discussion of this link to the Conclusion.

1.3. Contact, Interaction and Transmission. The probability of disease transmission during an act of unprotected sex between two individuals, is a function of the serostatus of the two individuals. In particular, no transmission is possible when the couple is *seroconcordant* (i.e. they share infection state), while it is possible when they are *serodiscordant*. Since infection is potentially present without any evident symptoms, the individuals must form beliefs about each other's (and indeed their own) infection state. The probability that the couple is serodiscordant (and thus that disease is potentially transmitted) is given by the *transmissive*

[2]Long-term relationships are assumed to be exogenously terminated if either partner shows symptoms of infection.

[3]In most of the existing literature, it is assumed that even though the health status of a given individual may evolve stochastically over time, there is no aggregate uncertainty. This implies that individuals' beliefs about aggregate disease prevalence are in fact always correct. An exception is the work by Chen (2009) in which individuals form beliefs about aggregate disease prevalence from observations of some (but not all) other individuals' health states.

contact rate

$$\kappa \equiv 1 - p_1 - p_4 \quad (1)$$
$$= p_2 + p_3 \quad (2)$$

where p_1 is the probability that neither individual is infected, p_4 is the probability that both are infected, p_2 is the probability that only the man is infected, while p_3 is the probability that only the woman is infected.[4] For transmission probability $\beta \in [0,1]$, this gives rise to the effective contact rate $\beta\kappa$, as long as the individuals decide to engage in unprotected sexual behavior. The equations (1)-(2) allow me to touch upon three issues that will become central to the analysis that follows.

First, note that the potential for transmission in a given coupling is directly related to the degree of serodiscordance and not necessarily to the individual infection probabilities of the individuals (i.e. to the marginal probabilities that the man and the woman are infected, p_M and p_W, respectively). Second, and related to the first point, is the observation that the correlation between the individuals' health states is a central measure to gauge potential disease transmission. Importantly though, the extent to which individuals' infection states are correlated is itself endogenous. There are two separate ways in which individuals' health can be dependent, namely through serosorting (i.e. through endogenous matching patterns) and through repeated interactions with the same partner(s). The former possibility is considered in Matthies and Toxvaerd (2015), while the latter possibility is considered with in the present paper. When two (or more) partners interact repeatedly over time, the infection probabilities will coevolve in interesting ways. The determination of this kind of coevolution is the main focus of this paper. Third, as equation (2) makes clear, to the extent that individuals are mainly interested in avoiding that they themselves become infected, their decisions will generally differ from those of a social planner who is concerned with the effective contact rate $\beta\kappa$.

Fourth, note that equations (1)-(2) are relevant *only* under unprotected sex, which may itself be the outcome of rational decision making. In other words, just because two individuals meet (or are matched), that does not imply that they will necessarily engage in unprotected sexual behavior. Philipson and Posner (1993) explicitly make the important distinction between *random matching* and *random interaction*. Most contributions to the literature, particularly those in the purely epidemiological tradition, treat the two notions interchangeably, by essentially assuming that two randomly matched individuals necessarily engage in unprotected behavior. This is the implicit assumption in most economic models of infectious disease, where disease incidence is taken to equal some analog of the transmissive contact rate $\beta\kappa$.

The benefit of such an approach is that the individuals' belief formation is particularly simple, as it is about aggregate proportions only. The drawback of this approach is that it makes a number of implicit assumptions that may be unintended and very special. In particular, it amounts to assuming that sexual relations are always (i) with random individuals, (ii) anonymous, (iii) one-off and (iv) unprotected. Thus tractability may come at a substantial price in terms of assumptions. The present paper thus has an advantage over most existing work, in that it makes none of these assumptions and may thus be seen as more realistic. On the other hand, it is less straightforward to aggregate the analysis to the population level in a tractable manner. Thus providing a satisfactory micro-foundation for the evolution of sexual decision making that relies on acceptable assumptions about the character of sexual

[4] To appreciate the parallels to the classical SI model with homogeneous mixing, consider picking two individuals independently from a large homogeneous population. Each such individual will be infected or susceptible with probabilities I and $(1-I)$ respectively, where $I \in [0,1]$ is the prevalence of infection in the population. The probability that such a match is serodiscordant is $I(1-I) + (1-I)I = 1 - (1-I)^2 - I^2$, which is the analog of the transmissive contact rate between two individuals in the present setting in the special case where $p_2 + p_4 = p_3 + p_4 = I$. This is the special case considered by Philipson and Posner (1993, p. 49).

interactions comes at the cost of complicating the extension to a population-level analysis.

Last, the fact that only serodiscordant relationships can transmit infection has lead to the notion that individuals may engage in so-called *serosorting*, i.e. positive assortative matching based on infection state (see e.g. Dow and Philipson, 1996, Philipson and Dow, 1998 and Greenwood et al., 2013). As a policy to reduce infection, serosorting remains controversial, as its effectiveness relies on the ability of individuals to assess their own and their potential partners' infection states (see e.g. Stoner et al., 2003 and Eaton et al., 2009) and on the absence of moral hazard (see Schroeder and Rojas, 2002).[5] Issues of risk attitudes, asymmetric information and matching are treated by Matthies and Toxvaerd (2015).

1.4. Overview of the Analysis. In this paper, I analyze the interdependent dynamics of beliefs over health states and the decision of whether to engage in unprotected sex, mainly in the context of a monogamous long-term relationship. In each period, each individual observes whether it or the partner shows any symptoms of the disease and updates beliefs in standard Bayesian fashion. Responding to these beliefs, the individuals then decide on whether to engage in transmissive (i.e. unprotected) behavior or protected behavior. While beliefs determine the sexual behavior of the couple, their decisions also influence their beliefs, since the choice between protected and unprotected sex determines whether an additional instance of exposure (and hence possible transmission of the pathogen) occurs.

The analysis focuses on four main themes. First, sexual decision making and belief formation has interesting interactions over time, which are time dependent. I show that when the couple engage in unprotected sex and observe no symptoms, each individual's infection probability can be non-monotone over time, first increasing and then decreasing. The reason for this phenomenon is that unprotected sex increases the correlation between the individuals' health states.

Second, I extend the analysis to concurrent partnerships. I show that when individuals have multiple concurrent relationships, the configuration of links between individuals matters. Furthermore, concurrency can induce multiple switches in sexual decisions even when no new symptoms of infection are observed.

Third, I show that when unprotected sex is voluntary, then the equilibrium amount of unprotected sex is lower than the amount chosen by a utilitarian social planner. The reasons is that when there is no coercion, the most at-risk individual is pivotal and may choose to block an unprotected relationship, even if engaging in it would maximize the aggregate welfare of the couple.

Last, I show that when unprotected sex is voluntary and individuals are altruistic (i.e. they take their partner's welfare into account), then there is more unprotected sex and infection than in the equilibrium with selfish preferences. I show that the outcomes under altruism nest those preferred by a Rawlsian and a utilitarian social planner, respectively. I furthermore show that the latter outcome can be achieved in a decentralized equilibrium, once transfers are introduced. I complete the analysis by considering the effects of self-abnegating preferences and the problem of determining the origins of the infection.

In Section 2, I set out the basic model. In Section 3, I take beliefs as given and derive the privately optimal decisions for the individuals. In Section 4, I analyze the dynamics of beliefs, given decisions by the individuals. In Section 5, I combine the results from the preceding sections to fully characterize the equilibrium dynamics of the model. In Section 6, I consider the possibility of concurrent relationships. In Section 7, I analyze welfare issues, the effects of transfers and extensions to altruistic and self-abnegating preferences, respectively. In Section

[5] That serosorting is imperfect is noted by Avery et al. (2001), who state that *"Frequently [...] risk imposition is not symmetric. A speeding or drunk driver does not drive along roads populated with like drivers, just as the person with a cold does not wander around in a world of cold sufferers"*.

8, I conclude and discuss several interesting extensions of the analysis. In the Appendix, I address the question of the most likely path and the identity of the index case (i.e. the problem of finding the most likely source of the infection). In the Appendix, I also outline several extensions of the basic model.

2. The Model

The basic building block of the model is a multi-period version of the framework introduced by Philipson and Posner (1993). Consider a monogamous couple (M, W) which in each period $t = 1, 2, ...$ engage in sexual relations.[6] While each individual in the couple derives utility from unprotected (or transmissive) sexual behavior, they also face the possibility that either or both can become infected with a sexually transmitted disease, such as HIV. At each encounter, the individuals must thus trade off the pleasure of unprotected behavior against the risks of being infected by the partner.

The infectious disease is possibly asymptomatic. For simplicity, I assume that an infected individual will show (publicly observable) symptoms with probability $\sigma \in [0, 1]$ in each period. Once infected, an individual remains infected (and infectious) in perpetuity. This is therefore a model in the *susceptible-infected* class (aka SI or as the *simple epidemic*). Similarly, if symptoms appear, the individual is known with certainty to be infected.

To formalize the tradeoffs facing the two individuals, let *net utility* from unprotected sex for individual $i = M, W$ be given by the reduced-form function[7]

$$U_i(r_{ij}) \equiv 1 - r_{ij}\beta\pi \tag{3}$$

where the transmission risk r_{ij} is the joint probability that i is susceptible and that the partner $j \neq i$ is infectious. The parameter $\beta \in [0, 1]$ is the infectiousness of the disease and $\pi > 0$ is the utility cost of becoming infected. In the special case where the infection states of the two individuals are independent, the transmission risk from individual j to individual i is given by $r_{ij} = p_j(1 - p_i)$. Note that for simplicity, it is assumed that unprotected sexual contact is the only possible source of infection.[8] This formulation of preferences is the simplest possible that makes the problem stationary.

The parameters (π, β, σ) are fixed throughout, while the probabilities r_{ij} and r_{ji} evolve endogenously over time. The parameters (β, π) enter directly into the individuals' utility functions, while the parameters (β, σ) influence how the individuals' beliefs evolve. To make the problem interesting, I impose the following:

Assumption 1. $\beta\pi > 1$.

This assumption ensures that $U_i(1) < 0$, i.e. that an individual who is known to be susceptible, would never choose to engage in unprotected sex with an individual who is known

[6] While the labels M and W are chosen to connote a heterosexual couple, this is only to fix ideas. The model is perfectly general and also applies to homosexual couples.

[7] The reduced form net utility from unprotected sex, can be micro founded as follows. Assume that the gross utility of a potentially susceptible individual $i = M, W$ from unprotected sex is $U_i^R \equiv A - B\alpha_R r_{ij}$ for some $A, B > 0$ and let the utility from protected sex be $U_i^S \equiv -C - B\alpha_S r_{ij}$, where $\alpha_S \in [0, \alpha_R]$ is the probability that protection fails and $\alpha_R > 0$ is the probability of transmission when no protection is used. This gives rise to the net utility function

$$\hat{U}_i(r_{ij}) \equiv U_i^R - U_i^S = (A + C) - (\alpha_R - \alpha_S)r_{ij}B$$

Normalizing so $U_i(r_{ij}) \equiv \hat{U}_i(r_{ij})/(A+C)$, defining $\pi \equiv B/(A+C)$ and letting $\beta \equiv \alpha_R - \alpha_S$ denote the increase in transmission probability when shifting from protected to unprotected sex, yields the reduced form function in the text. In the main analysis, I will assume that $\alpha_S = 0$, but will consider the case $\alpha_S > 0$ in Online Appendix C.

[8] Exogenous infection sources and imperfect protection is considered in Online Appendix C.

to be infected. Since $U_i(0) > 0$, the decision of whether to demand protected sex, is a simple function of the transmission risk that the individual faces. Throughout, I assume that the costs, benefits and transmission probabilities are symmetric across individuals. This assumption is for simplicity only and is inessential for the results in this paper. While unprotected sex potentially results in disease transmission, it is assumed that protected sex offers perfect protection.

Last, I assume that each individual behaves myopically, but only in the sense that it disregards the influence of its decisions on its future beliefs (although π can be interpreted as the discounted future disutility of becoming infected).[9] This type of assumption is commonplace in the literature (see e.g. Kremer, 1996, Philipson and Posner, 1993, Brito et al., 1991 and Kerwin, 2012) and very significantly simplifies the analysis. It amounts to the assumption that the individuals do not engage in active experimentation in their choices of unprotected versus protected behavior and that they in each period respond non-strategically to the current probabilities of becoming infected. While the simplification to the analysis is obviously an attractive feature, there are also good reasons to impose this assumption on purely behavioral grounds (much like the assertion that few would experiment while playing Russian roulette). This type of passive learning is considered in other contexts by e.g. Keller (2007) and references therein. It should be emphasized that under perfect state information (which is the de facto assumption maintained in most of the literature), optimal period-by-period prevention decisions are myopic. The reason is that the choice of whether to engage in protected sex today, does not commit an individual to engage in protected sex tomorrow.[10]

3. Privately Optimal Decisions

In this section, I characterize the privately optimal responses of the two individuals, taking as given their beliefs about transmission risks. The privately optimal choice for an individual $i = M, W$ in a given period t is given by

$$a_t^i(r_{ij}) = \begin{cases} 1 & if \quad U_i(r_{ij}) > 0 \\ 0 & if \quad U_i(r_{ij}) \leq 0 \end{cases} \tag{4}$$

where 1 stands for unprotected sex, while 0 stands for protected sex. I assume for now that there is no coercion and therefore the couple engages in protected sex if either (or both) prefer to protect.[11] The joint decision at time t is then given by

$$a_t(r_{ij}, r_{ji}) = \min\{a_t^i(r_{ij}), a_t^j(r_{ji})\} \tag{5}$$

Thus sex is protected whenever $\min\{a_t^i(r_{ij}), a_t^j(r_{ji})\} = 0$ and unprotected otherwise. Next, define the following indifference curves:

$$I_M \equiv \{(r_{MW}, r_{WM}) \in [0,1]^2 : U_M(r_{MW}) = 0\} \tag{6}$$

$$I_W \equiv \{(r_{MW}, r_{WM}) \in [0,1]^2 : U_W(r_{WM}) = 0\} \tag{7}$$

The curves (I_M, I_W), plotted in (p_M, p_W)-space in Figure 2, delineate the areas in which each individual desires unprotected and protected sex, respectively.

Figure 2 has a number of interesting features. Consider for example the man's decision problem and assume for simplicity that the health states of the individuals are independent. In

[9] When modeling preferences over health states when infection is (possibly) asymptomatic, one may think of health as a credence good.

[10] This should be contrasted to optimal vaccination and treatment decisions, which are necessarily forward-looking. See Toxvaerd (2014) for further discussion of this point.

[11] This is the assumption considered by Philipson and Posner (1993). In a later section, I analyze the problem when transfers between individuals are allowed and when the social norm is that unanimity is not needed to engage in unprotected sex.

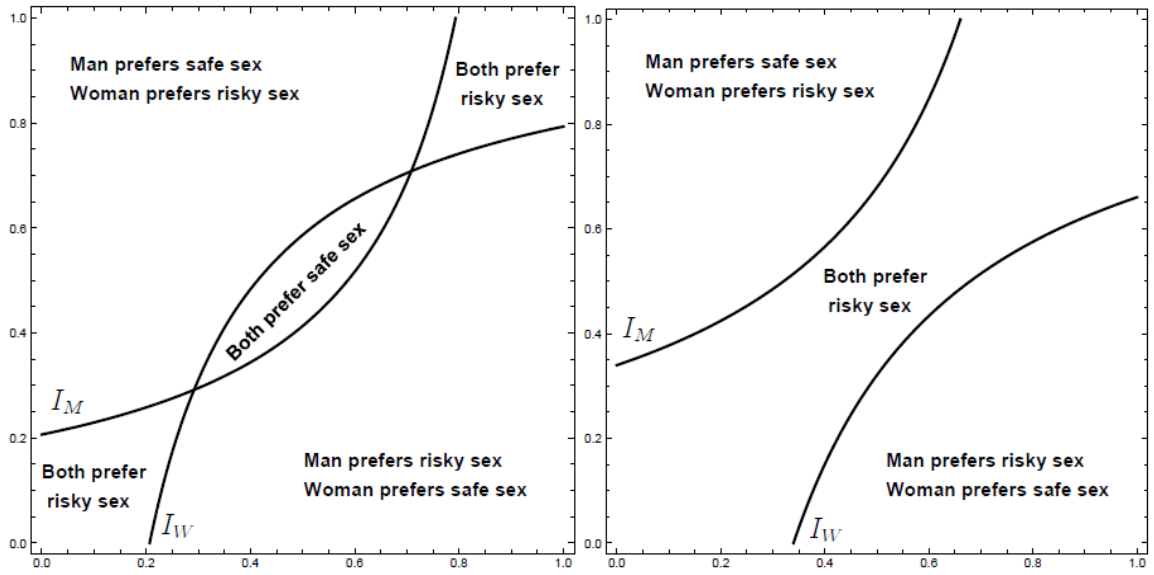

Figure 2: Optimal Decisions for Different Beliefs in (p_M, p_W)-Space. Contiguous scenario in left-hand side panel and non-contiguous scenario in right-hand side panel.

choosing whether to ask for protected or unprotected sex, the key statistic is the *effective transmission probability* $\beta(1-p_M)p_W$, rather than the probability p_W that the woman is infected. If p_W is very small, then the probability that the pathogen will be transmitted is accordingly small. Similarly, $(1-p_M)$ captures the man's expected susceptibility to infection. When this is low, which happens when the probability p_M that the man is already infected is high, then the probability of transmission is also low. The right-hand side panel of Figure 2 illustrates the case in which the area where both individuals prefer protected sex is empty. In conclusion, in the first case, when both partners have sufficiently high or sufficiently low probability of being infected, then the effective transmission probability is low and both partners will opt for unprotected sex. Otherwise, the couple will engage in protected sex. In the second case, there is a band around the 45°-line in which both individuals prefer unprotected sex.

Figure 2 gives a complete description of the individuals' decisions in (p_M, p_W)-space at a given point in time. For later use, I introduce the following notation:

$$\mathcal{S}_M \equiv \{(r_{MW}, r_{WM}) \in [0,1]^2 : U_M(r_{MW}) \leq 0\} \tag{8}$$
$$\mathcal{S}_W \equiv \{(r_{MW}, r_{WM}) \in [0,1]^2 : U_W(r_{WM}) \leq 0\} \tag{9}$$
$$\mathcal{S} \equiv \mathcal{S}_M \cup \mathcal{S}_W \tag{10}$$
$$\mathcal{R} \equiv \{(r_{MW}, r_{WM}) \in [0,1]^2 \backslash \mathcal{S}\} \tag{11}$$

Thus $\mathcal{S}_M$ and $\mathcal{S}_W$ are the sets of beliefs under which the man and the woman desire protected sex, respectively, and $\mathcal{S}$ is the set of beliefs under which the couple choose protected sex.[12] The set $\mathcal{R}$ is the set of beliefs under which both partners desire unprotected sex. The two scenarios illustrated in Figure 2 correspond to the scenarios in which the set $\mathcal{S}$ is contiguous and non-contiguous, respectively (and I shall henceforth refer to the two scenarios as such). Whether the relevant scenario is the contiguous or the non-contiguous one, depends on the parameters (β, π). The former scenario is more likely to obtain when the infectiousness β or the cost-benefit ratio π are sufficiently high (i.e. when the disease is very contagious or when

[12] Formally, $\mathcal{S}_M$ is the upper contour set of the man's indifference curve yielding zero net utility and $\mathcal{S}_W$ is similarly defined.

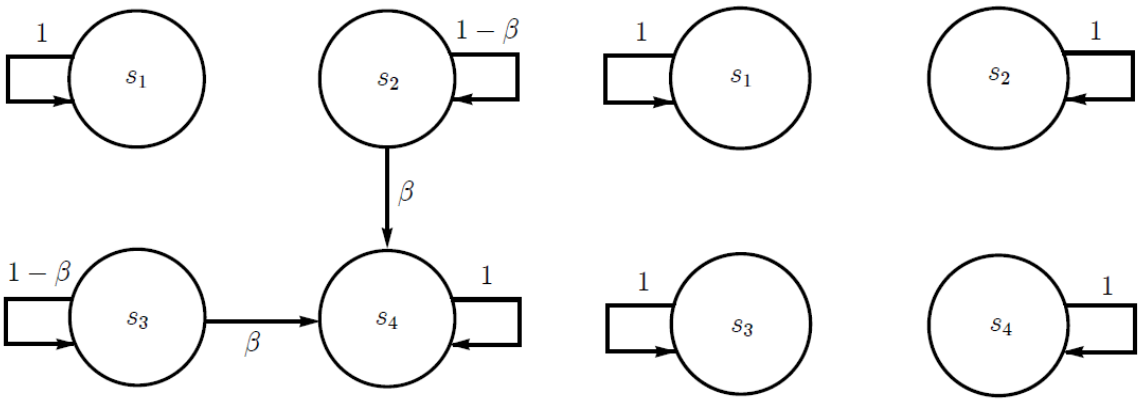

Figure 3: Markov State Transitions. Left-hand side shows case of unprotected sex; right-hand side shows case of protected sex.

the consequences of the disease are very severe).

In the general case in which the health states are dependent, the overall picture is the same as that for the independent case (which is of course a special case), but the exact location of the indifference curves is parameterized by the correlation coefficient. Since correlation will change endogenously over time, the location of the indifference curves will likewise change over time, which will be explored in detail in what follows. In particular, as correlation increases, the curves shift away from the diagonal so that the regime gradually shifts from the non-contiguous one to the contiguous one. In the limit, as the health states of the individuals become very positively correlated, the set $\mathcal{S}$ coincides with the space $[0,1]^2$.

4. Filtering and Belief Dynamics

In this section, I describe the evolution of beliefs for a given set of decisions by the individuals. While it may be tempting to characterize beliefs in terms of the infection probabilities of the individuals, i.e. in (p_M, p_W)-space, it turns out to be much more useful to work with the joint distribution over the possible health states of the *couple*. To this end, I start by introducing some notation.

Denote the set of possible states by

$$S = \{s_1, s_2, s_3, s_4\} \tag{12}$$

where

$$s_1 = (0,0), s_2 = (0,1), s_3 = (1,0), s_4 = (1,1) \tag{13}$$

The first coordinate of a state is the infection state of M, while the second coordinate is the infection state of W; here 1 stands for infected and 0 stands for susceptible. Let the random state at time $t = 0, 1, ...$ be denoted by $X_t \in S$ and denote a realization by x_t.

Let the prior distribution over states be given by $P(X_0) = (p_1, p_2, p_3, p_4)$. The state transition matrix is given by

$$T(a_{t-1}) \equiv \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1-\beta a_{t-1} & 0 & \beta a_{t-1} \\ 0 & 0 & 1-\beta a_{t-1} & \beta a_{t-1} \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{14}$$

where entry τ_{ij} denotes the probability that the system transitions from state i to state j and a_t is the joint decision on the type of sex at time t. The state transition matrix under protected

sex is simply the identity matrix.

It is immediately clear that under (perpetual) unprotected sex, the states $\{s_1, s_4\}$ are absorbing, while the states $\{s_2, s_3\}$ are transient. Under (perpetual) protected sex however, all states are absorbing. Note also that under unprotected sex, transitions from the transient states $\{s_2, s_3\}$ to s_4 are Bernoulli trials with success parameter β and thus the time to infection (of the remaining susceptible individual) is geometrically distributed with mean time $1/\beta$.[13] Similarly, the expected time to transition from asymptomatic to symptomatic for an infected individual, is geometrically distributed with mean time $1/\sigma$.

It should be noted that the Markov chain characterized by the transition matrix $T(a_{t-1})$ is not irreducible and that it contains two closed classes, namely $\{s_1\}$ and $\{s_4\}$. As a consequence, the Markov chain does not have a unique invariant distribution, but three. For a fixed joint decision $a \in \{0, 1\}$, these are given by

$$(p_1, (1-a)p_2, (1-a)p_3, a(p_2+p_3)+p_4) \tag{15}$$

and the degenerate distributions $(1, 0, 0, 0)$ and $(0, 0, 0, 1)$.

Figure 3 illustrates the states and transitions for the two stationary Markov chains induced by perpetual unprotected and protected sex, respectively. Note that for clarity, only state transitions that have positive probability are indicated.

Although Figure 3 gives a complete description of the state transitions under the two possible kinds of sex, it is not helpful in describing the decision making process of the individuals, since the states are not always observed directly. While there is an underlying process that determines the evolution of the states over time (given the individuals' decisions), all decision making is based on an observed process of possible symptoms, whose states depend on the underlying process. Inference problems with this formal structure are known as Hidden Markov Models (HMMs) and are commonplace in the literatures on speech processing and on bioinformatics.[14] An HMM is formally composed of a transition model and a sensor model. The former is simply the underlying Markov chain described above. The latter is given by a message space (or *alphabet*) and a set of (state dependent) emission probabilities. For the current context, let the message space be given by

$$R = \{r_1, r_2, r_3, r_4\} \tag{16}$$

where

$$r_1 = (0, 0), r_2 = (0, 1), r_3 = (1, 0), r_4 = (1, 1) \tag{17}$$

The first coordinate of a message is the symptom state of M, while the second coordinate is the symptom state of W; here 1 stands for symptomatic and 0 stands for asymptomatic. Let the random message (or symptom observations) at time $t = 0, 1, ...$ be denoted by $Y_t \in R$ and denote a realization by y_t. Last, the emission probabilities are given by the matrix

$$Q \equiv \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1-\sigma & \sigma & 0 & 0 \\ 1-\sigma & 0 & \sigma & 0 \\ (1-\sigma)^2 & (1-\sigma)\sigma & \sigma(1-\sigma) & \sigma^2 \end{pmatrix} \tag{18}$$

where entry q_{ij} denotes the probability that the message y_j is observed in state s_i.

To summarize, while utilities depend directly on the state sequence $\{X_t\}_{t=1}^{\infty}$, while decisions are made based on the symptom sequence $\{Y_t\}_{t=1}^{\infty}$. Figure 4 illustrates the state transitions

[13] Pinkerton and Abramson (1993) consider such a model, but disregard symptoms and beliefs.

[14] See e.g. Rabiner (1989), Russell and Norvig (2003) or Fristedt et al. (2007) for more details.

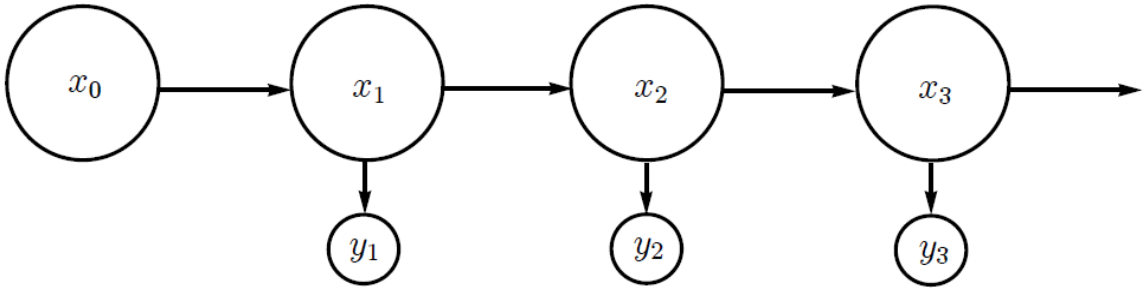

Figure 4: Unobservable State Transitions and Symptom Observations.

and the sequence of observed symptoms. The important simplification of the HMM structure, is that the state at any given time t, depends only on the state and decisions at time $t-1$ (and not on states or symptom observations at previous times), while the symptom observations at time t depend only on the state at time t (and not on states at times before time t).

The evolution of beliefs can be calculated recursively by the so-called *forward equation*, given by

$$P(X_t|y_{1:t}) = \frac{P(y_t|X_t)\sum_{x_{t-1}} P(X_t|x_{t-1})P(x_{t-1}|y_{1:t-1})}{\sum_{X_t} P(y_t|X_t)P(X_t|y_{1:t-1})} \tag{19a}$$

$$= \nu O_t T'(a_{t-1})P(x_{t-1}|y_{1:t-1}) \tag{19b}$$

for $t = 1, 2, ...$, where $y_{1:t} \equiv (y_1, ..., y_t)$ denotes the observed messages in periods 1 through t and

$$\nu \equiv \frac{1}{\sum_{X_t} P(y_t|X_t)P(X_t|y_{1:t-1})} \tag{20}$$

is a normalizing constant. The observation matrix O_t is a 4×4 diagonal matrix whose diagonal elements are the elements of the relevant column of the matrix Q (i.e. the column that is relevant given the current symptom observations) and which has zeros everywhere else.

To understand the characterization of belief dynamics below, it is instructive to interpret the forward equation. The formation of beliefs can be decomposed into two steps, namely *prediction* and *updating*.

The prediction step is simply the propagation of previous period's beliefs, according to the underlying Markov chain. Specifically, this step is given by

$$P(X_t|y_{1:t-1}) = \sum_{x_{t-1}} P(X_t|x_{t-1})P(x_{t-1}|y_{1:t-1}) \tag{21}$$

Next, the updating step takes this prediction and updates it according to Bayes' rule, in light of the new symptom observations. Specifically, this step is given by

$$P(X_t|y_{1:t}) = \frac{P(y_t|X_t)P(X_t|y_{1:t-1})}{\sum_{X_t} P(y_t|X_t)P(X_t|y_{1:t-1})} \tag{22}$$

Combining the prediction step and the updating step, one readily obtains the forward equation. It is sometimes more informative to consider the posterior state probabilities individually; these are therefore listed in Online Appendix A.

To build intuition, it is useful to consider two special cases of the forward equation that are readily understood, namely the case where $\beta = 0$ and $\sigma \in (0, 1)$ and the case $\sigma = 0$ and $\beta \in (0, 1)$. In the former case, the prediction step is based on the identity matrix and hence

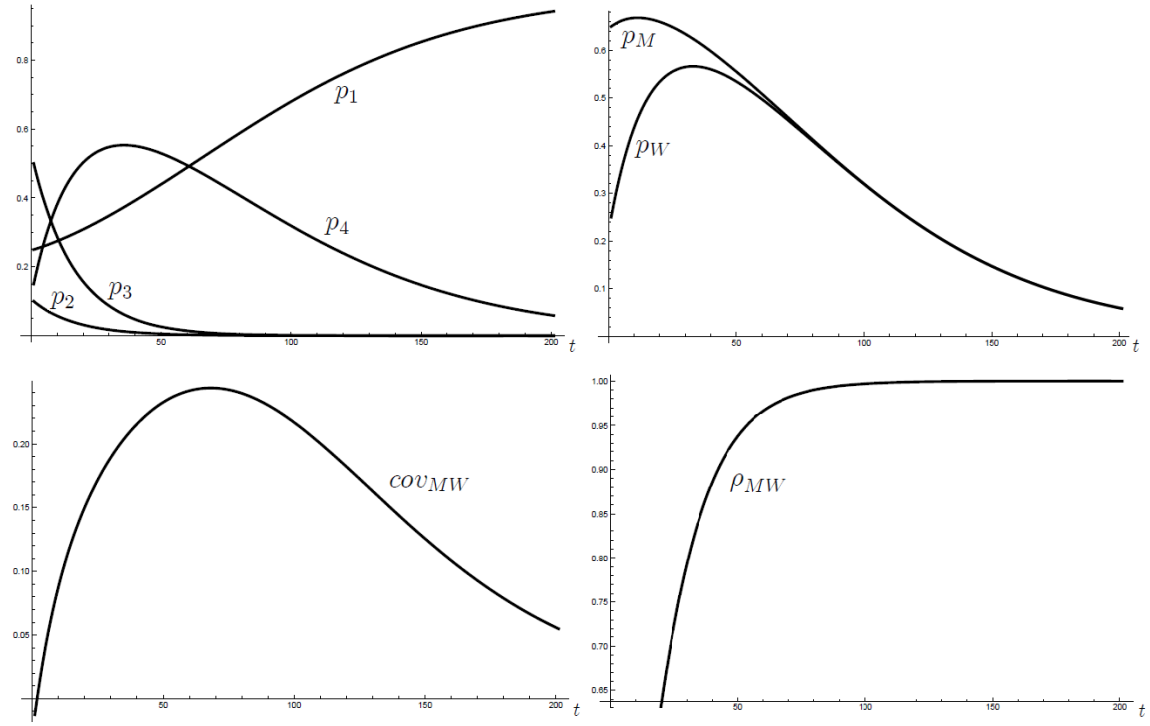

Figure 5: State Probabilities, Marginals, Covariance and Correlation Without Symptoms.

the state does not change over time. Thus, the forward equation reduces to Bayes' rule, which gives the posterior distribution over states (between which there are no transitions), given the history of observed signals. In the latter case, the prediction step is based on the transition matrix T and no symptom observations are ever made or used in updating. To update beliefs, one simply propagates the Markov process blindly, i.e. without the use of observations. In other words, beliefs are based entirely on the transition model and no use is made of the sensor model. Away from these extreme cases, when $\beta, \sigma \in (0,1)$, the forward equation nests these two different forecasting rules.

As an alternative to the Markov chain representation of beliefs, one can consider the health state of each individual as a Bernoulli random variable and hence describe beliefs as a partially observable bivariate Bernoulli process. Denote the health state of $i = M, W$ by $H_i \in \{0,1\}$, with typical element h_i. It is straightforward to show that

$$cov(H_M, H_W) = p_1 p_4 - p_2 p_3 \equiv cov_{MW} \tag{23}$$

The covariance takes values in the interval $[-1/4, 1/4]$. In turn, the correlation is given by

$$\rho_{MW} = \frac{p_1 p_4 - p_2 p_3}{\sqrt{(p_1 + p_2)(p_1 + p_3)(p_4 + p_2)(p_4 + p_3)}} \tag{24}$$

which takes values in the interval $[-1, 1]$. In the next subsection, I will explain how beliefs can be characterized in (p_M, p_W)-space. For this purpose, it is essential to appreciate exactly how the health states of the two individuals covary and how the covariance changes over time.

Figure 5 illustrates the evolution of beliefs under unprotected sex when no symptoms are observed, for the special case $\beta = 0.06$ and $\sigma = 0.025$. This is the parameterization used by Greenwood et al. (2013) in their calibration. The figure indicates the interesting result that for empirically plausible values of the parameters, beliefs can be non-monotone over time, even

without the individuals observing symptoms or switching away from unprotected sex. The reason for this result will be made clear shortly.

The upper left-hand side panel of Figure 5 shows the evolution over time of the state probabilities (p_1, p_2, p_3, p_4). I have chosen different initial conditions for p_2 and p_3, so that their curves do not overlap. There are several interesting features worth mentioning. First, the probabilities that only one but not the other individual is infected, i.e. p_2 and p_3, are strictly decreasing over time, approaching zero. Second, the probability p_1 that neither is infected, increases over time, approaching one. Last, the probability p_4 that both are infected, is hump-shaped. The upper right-hand side panel shows the marginal probabilities that the two individuals are infected, i.e. p_M and p_W, respectively. Again, these probabilities are hump shaped, which follows from the fact that $p_M = p_3 + p_4$ and $p_W = p_2 + p_4$. In other words, the non-monotonicity in the marginal infection probabilities stems from the behavior in the state probability p_4 that both individuals are infected and not from the probabilities p_2 and p_3 that either individual is infected on its own. The lower left-hand side panel shows the covariance between the health states of the two individuals, which first increases and then decreases. The lower right-hand side panel shows the correlation between the health states of the two individuals. As can be seen from the figure, the correlation strictly increases over time, tending to one. It is also the case that over time, if no symptoms are observed, then the individuals' health states become increasingly correlated even under protected sex. But the rate at which the correlation increases in this case, is lower than the rate under unprotected sex.

At first blush, it is perhaps counter-intuitive that beliefs can be increasing over time even though no symptoms have been observed. It is therefore appropriate to emphasize the reasons behind this finding, which are found in the prediction and updating steps described above. When calculating the probability of the individuals being asymptomatically infected, note that the prediction step is based on $T(a_{t-1})$, which is a function of whether unprotected or protected sex is engaged in. During the prediction step, the "interim posterior", i.e. the probability assigned to the event that the individual is infected, but before additional information has been observed, is weakly higher than the prior probability assigned to that event. In case of protected sex, it is in fact unchanged, while under unprotected sex, it is weakly higher because either the individual is already infected or he or she now becomes infected with positive probability.

Turning to the updating step, consider first the case in which the individual had protected sex. In this case, the fact that no symptoms were observed means that the posterior probability of being asymptomatically infected must be lower than the prior. This is because there has been no additional exposure, but the individual has received an encouraging signal that he or she is not infected. Things are more complicated if the individual engaged in unprotected sex. In this case, the interim posterior may be higher than the prior, but the fact that no symptoms have been observed, acts to decrease the probability assigned to asymptomatic infection, *from the benchmark of the interim posterior*. Thus the posterior belief weighs the increase in probability assigned to asymptomatic infection caused by unprotected sex, against the decrease caused by the absence of symptoms. The net effect depends on the prior beliefs over the two individuals' health states and on the transmission and precision parameters (β, σ).

An alternative way to understand the non-monotonicity, is to note that there can never be any transitions into states $\{s_1, s_2, s_3\}$. This means in particular that as long as no symptoms are observed, the probability of these states are monotone. As a lack of symptoms is (imperfect) evidence of the absence of infection, the probability of state s_1 is monotone increasing while the probabilities of states s_2 and s_3 are monotone decreasing. Turning to state s_4, things are somewhat different, because there can be transitions into this state from states $\{s_2, s_3\}$. While it remains true that a lack of symptoms is evidence of an absence of infection, this must be traded off against the possibility of a transition to infection. These ideas are most clearly seen

in the expressions listed in Online Appendix A.

In short, at early stages of the relationship, additional exposure increases the probabilities that the individuals are infected, despite no symptoms being observed. As time passes, the health states become increasingly correlated, leading to infection probabilities eventually declining (as long as no symptoms are observed).

If the individuals' health states are perfectly correlated, beliefs must be on the diagonal in (p_M, p_W)-space and the only remaining uncertainty is whether they are both infected or both susceptible. In fact, the sum $p_2 + p_3$ is the probability assigned to the couple being serodiscordant, while the sum $p_1 + p_4$ is the probability that they are seroconcordant. When correlation is perfect, each individual gets two signals about its own health state in each period (and hence learns very fast), irrespective of the choice over protection. But they can clearly also be on the diagonal without being correlated.

4.1. Tracking Beliefs in (p_M, p_W)-Space. When characterizing the privately optimal decisions of the individuals, I did so in (p_M, p_W)-space, while I moved to (p_1, p_2, p_3, p_4)-space in order to characterize beliefs. In this section, I will recast the characterization of beliefs in (p_M, p_W)-space and draw out some of the complications involved in doing so.

Suppose that both partners are asymptomatic and engage in continual unprotected sex. Also, start by considering the special case in which the two individuals' health states are independent. It is useful to fix ideas by considering the inference problem at time $t = 1$. In this period, the prior distribution over states is given by

$$(p_1, p_2, p_3, p_4) = ((1 - p_M)(1 - p_W), (1 - p_M)p_W, p_M(1 - p_W), p_M p_W) \tag{25}$$

That is, before the two individuals have engaged in unprotected behavior, their health states are completely independent. But once they have engaged in unprotected sex (but are still asymptomatic), what can one say about their posterior beliefs (p'_M, p'_W) compared to their priors (p_M, p_W)? It turns out that there is a simple way of characterizing the change in beliefs after this first unprotected encounter. Specifically, for any pair (β, σ), there exists a pair of increasing functions J_M and J_W in (p_M, p_W)-space defined by

$$J_M \equiv \{(r_{MW}, r_{WM}) \in [0,1]^2 : p_M = p'_M\} \tag{26}$$
$$J_W \equiv \{(r_{MW}, r_{WM}) \in [0,1]^2 : p_W = p'_W\} \tag{27}$$

where the posterior infection probabilities, conditional on both individuals being asymptomatic, are given by

$$p'_M \equiv P(X_t = x_3 | y_t = r_1) + P(X_t = x_4 | y_t = r_1) \tag{28}$$
$$p'_W \equiv P(X_t = x_2 | y_t = r_1) + P(X_t = x_4 | y_t = r_1) \tag{29}$$

Along these iso-belief curves, the probabilities that the individuals are infected but asymptomatic remain unchanged from one period to the next.

The curves (J_M, J_W), plotted in (p_M, p_W)-space in Figure 6, delineate the areas in which each posterior infection probability is higher (respectively lower) than the prior probability, as a function of the prior probability assigned to the partner being infected (but asymptomatic). The functional forms of the iso-belief curves J_M and J_W given by the two equations

$$p_W = \frac{p_M \sigma(1 + \beta(1 - \sigma))}{p_M \sigma^2 + \beta(1 - \sigma)(1 - \sigma + 2p_M \sigma)} \tag{30}$$
$$p_W = \frac{p_M \beta(1 - \sigma)^2}{\sigma(1 + \beta - 2p_M \beta - p_M \sigma - \beta\sigma + 2p_M \beta\sigma)} \tag{31}$$

While the functions (J_M, J_W) are both increasing in (p_M, p_W)-space, the magnitude of (β, σ) determine the location of the functions. Depending on parameters, there are three cases to consider as follows:

Case A. $J_M \leq J_W$ for all p_M.

Case B. $J_M \geq J_W$ for all p_M.

Case C. $J_M \geq J_W$ for $p_M \leq \widetilde{p}$ and $J_M \leq J_W$ for $p_M \geq \widetilde{p}$ for some $\widetilde{p} \in [0, 1]$.

In Case C, the threshold $\widetilde{p}$ is found as the interior element of the set $\{J_M \cap J_W\}$. Specifically, it is given by

$$\widetilde{p} \equiv \frac{2\beta\sigma^2 - 3\beta\sigma + \beta - \sigma}{2\beta\sigma^2 - \sigma^2 - 2\beta\sigma} \tag{32}$$

The parameter constellations that give rise to the three cases are illustrated in (β, σ)-space in Figure 7. Case A obtains for $\beta >> \sigma$, while Case B obtains for $\sigma > \beta$. Case C obtains for $\beta > \sigma$ but with β not too much larger than σ. Cases A and B are readily interpreted. When $\beta >> \sigma$, the probability of having been infected through unprotected sex, outweighs the fact that no symptoms have been observed. Similarly, when $\sigma > \beta$, the fact that there have been no symptoms, weighs more than the fact that the couple have engaged in unprotected sex and thus exposed themselves to the possibility of infection. Case C constitutes the hybrid case, which is similar to Case A for high infection probabilities and similar to Case B for low infection probabilities. It is interesting to note that the parameter estimates used by Greenwood et al. (2013), belong to Case A in which beliefs are non-monotone.

Next, turn to some period $t > 1$. The inference problem faced in any such period is qualitatively different from the one faced at $t = 1$, even controlling for the prior infection probabilities. The reason is that over time, the two individuals' health states become increasingly correlated. That is, after the first period with unprotected sex, it is no longer legitimate to treat them as independent. While the covariance may be non-monotone, the correlation is always monotone (increasing) in the number of unprotected contacts.

It turns out that when the correlation between the individuals' health states is increased, the iso-belief curves J_M and J_W shift in a north-west and south-east direction, respectively. This means that irrespective of the initial regime, over time the movement is always in the direction of Case A, i.e. the case in which the beliefs are decreasing when no symptoms are observed, even when individuals engage in unprotected sex.[15] In a sense, increased correlation works in a similar way as an increase in the signal precision σ. This is not a coincidence. When the health states are completely independent, each symptom observation is informative only about the health state of the given individual. As health states become more correlated, each signal becomes more informative about both health states and the individuals therefore receive more precise information in each period. In the limit, as the health states become perfectly correlated, the individuals receive two signals about the (common) health state each period. This corresponds qualitatively to an increase in signal precision σ.

Having derived both the indifference curves I_M and I_W and the iso-belief curves J_M and J_W, it is tempting to seek to determine the equilibrium dynamics by superimposing the curves and tracking beliefs and decisions in (p_M, p_W)-space. The different possible cases are illustrated in Figure 16 in Online Appendix D, which is a snapshot at a given point in time and assumes that the health states of the individuals are independent (i.e. they are independently drawn and that the individuals are at the start of their relationship). It is clear from the figure that

[15] For completeness, it should be noted that when starting in Case B, increasing the correlation moves the iso-belief curves towards something resembling Case C, but where the lower intersection may be at positive probabilities rather than at the origin, as is the case when the health states are independent.

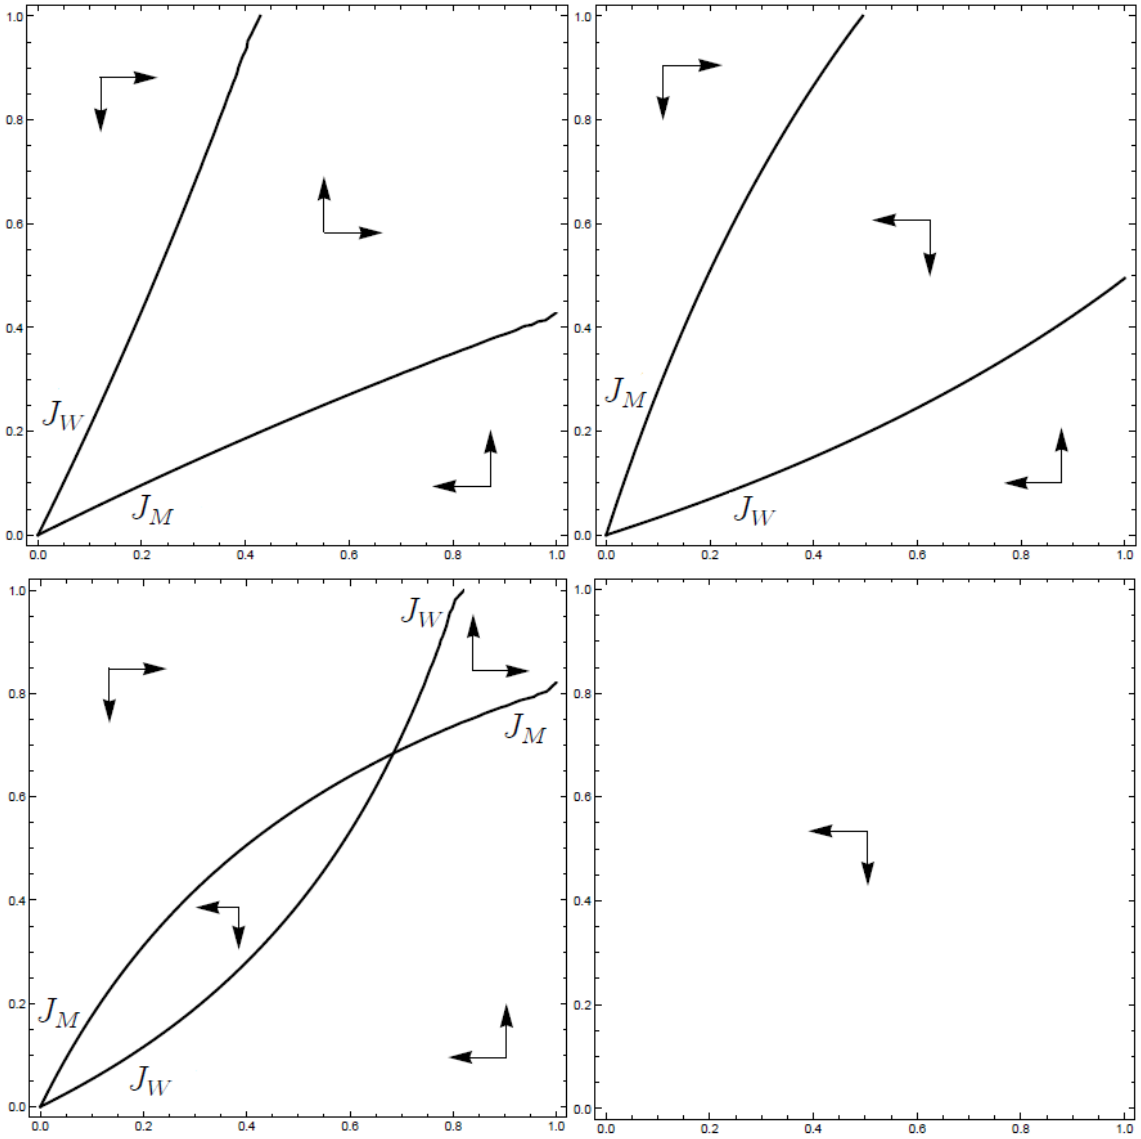

Figure 6: Belief Dynamics Under Risky and Safe Sex in (p_M, p_W)-Space. Upper left-hand side panel is Case A, where $J_M \leq J_W$. Upper right-hand side panel is Case B, where $J_M \geq J_W$. Lower left-hand side panel is Case C, where $J_M \geq J_W$ for $p_M \leq \widetilde{p}$ and $J_M \leq J_W$ for $p_M \geq \widetilde{p}$. Lower right-hand side panel shows belief dynamics under safe sex.

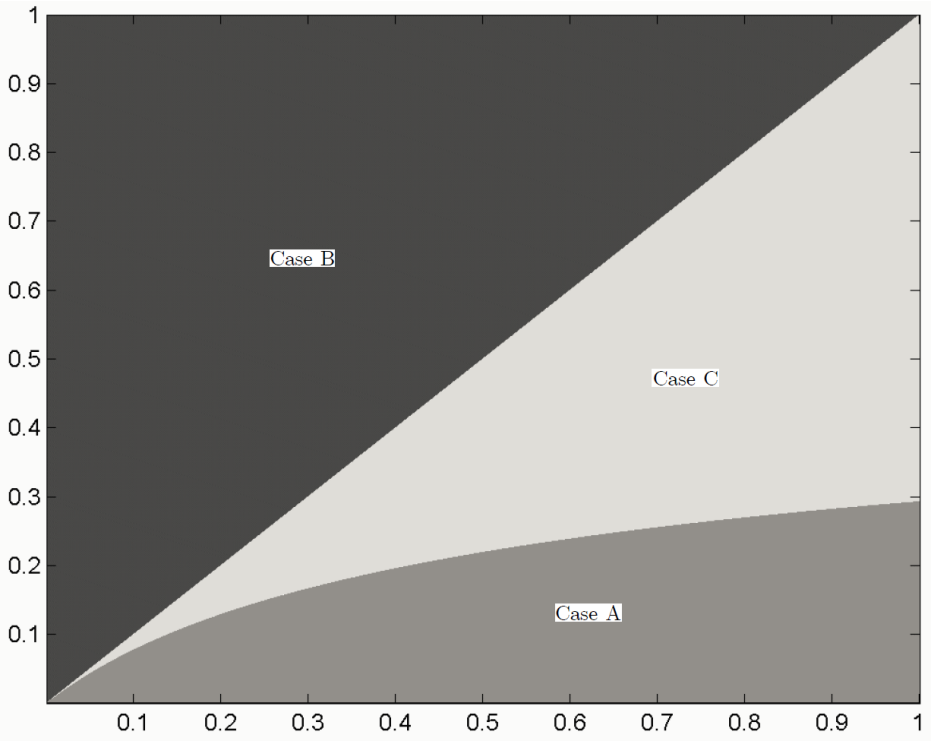

Figure 7: Parameter Constellations for Cases A, B and C in (β, σ)-Space.

there is a great many cases to consider, even in the first period when the health states of the individuals are independent. As time passes and the health states become increasingly correlated, all the curves move around at different speeds. This fact makes it very complicated to work directly in (p_M, p_W)-space.

5. Equilibrium Decisions and Paths

In the previous sections, I have determined how decisions are made for given beliefs and characterized how beliefs evolve for given decisions. I now combine these results to characterize the equilibrium dynamics of decisions and beliefs.

Define the critical threshold $\bar{p} \in [0, 1]$ at which $U_i(\bar{p}) = 0$. It is given by

$$\bar{p} \equiv (\beta\pi)^{-1} < 1 \tag{33}$$

In what follows, it is important to note that $r_{MW} = p_2$ and $r_{WM} = p_3$. That is, the transmission risk faced by an individual, equals the joint probability that the individual is susceptible and that the partner is infected. I next consider the individuals' privately optimal decision rules. They can be expressed as

$$a^M(r_{MW}) = \begin{cases} 1 & if \quad p_2 < \bar{p} \\ 0 & if \quad p_2 \geq \bar{p} \end{cases} \tag{34}$$

$$a^W(r_{WM}) = \begin{cases} 1 & if \quad p_3 < \bar{p} \\ 0 & if \quad p_3 \geq \bar{p} \end{cases} \tag{35}$$

When the partner of an individual is known to be infected, the decision rules specialize to

$$a^M(1-p_M) = \begin{cases} 1 & if \quad p_M > 1-\bar{p} \\ 0 & if \quad p_M \leq 1-\bar{p} \end{cases} \tag{36}$$

$$a^W(1-p_W) = \begin{cases} 1 & if \quad p_W > 1-\bar{p} \\ 0 & if \quad p_W \leq 1-\bar{p} \end{cases} \tag{37}$$

The probability of being in states s_2 and s_3, respectively, decrease over time when no symptoms are observed.[16] These are the state probabilities that matter for the individuals' decision making. Thus, if the individuals start by engaging in unprotected sex, they will never switch away from doing so, unless symptoms appear in one or the other individual (but not if symptoms are simultaneously observed in both).

Because it has been assumed that there is no coercion, it is always the decision of the most at-risk individual that determines the type of sex that the couple engages in. To see this, note that $U_i(r_{ij}) > U_j(r_{ji})$ if an only if $r_{ji} \geq r_{ij}$. In other words, the individual facing the highest transmission risk, will have the lowest propensity to engage in unprotected behavior. Suppose that $r_{MW} \geq r_{WM}$. It then follows from (28) and (29) that $p_M \leq p_W$ and furthermore, as long as no symptoms are observed, the posterior beliefs will be such that $p'_M \leq p'_W$. Therefore, it is always the lowest probability individual that is pivotal and hence one can trace the couple's joint decisions by determining the decisions of that individual.[17] The observation that the most at-risk individual is always pivotal for the decision to engage in unprotected sex, means that the game can be treated as a single-player Markov decision process.
I can now state the following characterization of equilibrium behavior:

Theorem 1. *The equilibrium decisions and paths are as follows:*
(i) For all $t = 1, 2, ...$ such that $y_t = r_1$,

$$a^*_t = \begin{cases} 1 & if \quad \max\{p_2, p_3\} < \bar{p} \\ 0 & if \quad \max\{p_2, p_3\} \geq \bar{p} \end{cases} \tag{38}$$

Along such a path, there can be at most one switch in actions from $a^*_t = 0$ to $a^*_t = 1$, which follows from the observation that p_2 and p_3 are decreasing over time as long as $y_t = r_1$ is observed. Along such a path, beliefs converge to $(1, 0, 0, 0)$.
(ii) For some $s = 1, 2, ...$ such that $y_s = r_2, r_3$,

$$a^*_t = \begin{cases} 1 & if \quad \min\{p_M, p_W\} > 1-\bar{p} \\ 0 & if \quad \min\{p_M, p_W\} \leq 1-\bar{p} \end{cases} \tag{39}$$

for all $t > s$. As long as no further symptoms are observed, there are no switches in actions and beliefs converge to either $(0, 1, 0, 0)$ or $(0, 0, 1, 0)$ (depending on which individual showed symptoms at time s).
(iii) For some $s = 1, 2, ...$ such that $y_s = r_4$, $a^*_t = 1$ for all $t > s$ and beliefs remain at $(0, 0, 0, 1)$.

[16] In general, the joint probabilities (r_{ij}, r_{ji}) determine the transmission probabilities between two individuals. Under a long-term monogamous relationship, these are necessarily decreasing in time as long as no symptoms are observed. The reason is that no individual can be the source of their own infection. This means that as the probability that the partner was infected from the outset decreases, so must the probability that this partner has transmitted his infection during unprotected sex. When one allows for concurrent relationships, this is no longer the case.

[17] This feature follows from the symmetry that the parameter β is common to the individuals. If one party is intrinsically more susceptible than the other (as is the case with women versus men), then the individual with the highest infection probability, need not be the individual who is least at risk of infection from unprotected sex.

In words, the results can be understood as follows. In situations in which both individuals display symptoms, there is no need for any protection as there can be no disease transmission between seroconcordant individuals. The beliefs will in this case remain at $(0,0,0,1)$ indefinitely. On the other hand, when no symptoms have been observed, then the joint decision will follow that of the most at-risk individual. If this individual is sufficiently confident that the transmission risk is low (which will be the case after a sufficiently long period without symptoms), then he will agree to engage in unprotected sex. Otherwise, he will insist on protection until beliefs about transmission risks have decreased sufficiently to justify a switch to unprotected sex. In either case, if no further symptoms are observed, beliefs converge to $(1,0,0,0)$.

Last, consider the problem of a man whose partner has just become symptomatic. For this individual, the problem is particularly simple. Since the woman is assumed to be symptomatic, as long as the man is asymptomatic, any remaining dynamics take place on the upper edge in (p_M, p_W)-space. The only question for the man is therefore to determine the location of his posterior $P(X_t = s_4|y_t = r_2)$ vis-à-vis the iso-belief curve J_M. For posteriors $P(X_t = s_4|y_t = r_2) > 1 - \overline{p}$, beliefs are in the set $\mathcal{R}$ and the man will therefore find it optimal to engage in unprotected sex from then on, even without having himself shown any symptoms of infection. As argued above, under continual unprotected sex with an infected partner, the probability of being infected increases over time, tending to one. Thus in the limit, both individuals are infected and the system is absorbed in state s_4. For posteriors $P(X_t = s_4|y_t = r_2) \leq 1 - \overline{p}$, beliefs are in the set $\mathcal{S}_M$ and therefore the man will find it optimal to insist on protected sex from that point onward. This means that he will no longer be exposed to new infection. As a consequence, as long as he remains asymptomatic, his probability of being infected tends to zero. If the man is indeed infected, then he will eventually become symptomatic too and the beliefs will jump to the absorbing state $(0,0,0,1)$. If on the other hand he is in fact not infected, then over time, the beliefs will tend to the point $(0,1,0,0)$ and the couple will continue to engage in protected sex only. From this case, it is immediately clear that there is path dependence, in the sense that the state of beliefs at the time that symptoms appear, determine the subsequent evolution of decisions and beliefs. In state s_4, the only possible steady state is $(0,0,0,1)$. In states s_2 and s_3, $(0,0,0,1)$ is a feasible stable steady state, as are the additional points $(0,1,0,0)$ and $(0,0,1,0)$, respectively.

In this model, learning is always complete, in the sense that beliefs always converge to the true state of the world (which is itself endogenously determined). The reason is that any infected individual will eventually display symptoms (and beliefs will reflect this perfectly). The asymptomatic partner of a symptomatic individual will either engage in unprotected sex or in protected sex, in perpetuity. In the former case, the asymptomatic partner will become infected and symptomatic in the limit. In the latter case, the asymptomatic partner will eventually show symptoms (in case he or she is in fact infected) or never show symptoms, in which case the probability that this individual is infected will approach zero.

For comparison, the long-term beliefs in the absence of the observation process are given as a corollary:

Corollary 2. *When no symptoms are available, the equilibrium decisions and paths are as follows:*
(i) If $\max\{p_2, p_3\} < \bar{p}$, *then* $a_t^* = 1$ *for all* $t = 1, 2, ...$ *and beliefs converge to* $(p_1, 0, 0, p_2+p_3+p_4)$.
(ii) If $\max\{p_2, p_3\} \geq \bar{p}$, *then* $a_t^* = 0$ *for all* $t = 1, 2, ...$ *and beliefs remain at* (p_1, p_2, p_3, p_4).

This result follows from Theorem 1 and the invariant distribution in (15).

5.1. A Brief Return to (p_M, p_W)-Space. It is worthwhile retuning to (p_M, p_W)-space to appreciate the complications involved in working there, rather than in (p_1, p_2, p_3, p_4)-space.

Consider a parameter constellation in Case A in the contiguous scenario, illustrated in the left-hand side panel of Figure 8. Suppose initial beliefs are at a point a on the diagonal. With these beliefs, the individuals will choose to engage in unprotected sex and if they observe no symptoms, the beliefs will increase in the subsequent period(s). As the individuals engage in unprotected sex, their health states become increasingly correlated, causing two separate effects. On the one hand, the iso-belief curves start shifting towards Case C and eventually towards Case B. This means that if no symptoms are observed, beliefs will decrease towards a point like b and then onwards to a point like c. On the other hand, the increase in correlation causes a shift in the indifference curves away from the diagonal. In other words, the scenario moves from the contiguous one to the non-contiguous one. The important question that arises is whether, starting at a point $a \in \mathcal{R}$, beliefs can converge to a point $b \in \mathcal{S}$ and on to a point $c \in \mathcal{R}$. If this were possible, it would imply that on an equilibrium path on which no symptoms are observed, individuals would switch from unprotected to protected sex and then switch back to unprotected sex. The answer to this question depends on the strength of the two effects, i.e. on the relative speeds at which the indifference curves and the iso-belief curves move as correlation is increased. Using the Markov chain representation of the health states, this question is readily settled. The answer is that such multiple switches cannot happen in equilibrium. This would have been extremely difficult to establish based on an analysis in (p_M, p_W)-space.

As noted earlier, when starting in Case A, initially the marginal infection probabilities of the individuals increase as the individuals engage in unprotected sex. That is, even though no symptoms have been observed in either individual, the initial probabilities assigned to the risk of infection are sufficient to increase posterior beliefs of being infected. Over time though, as no symptoms are observed, it becomes increasingly likely that neither individual is in fact infected. At this stage, posteriors start decreasing despite there being additional exposure. This effect is captured in the increase in correlation depicted in Figure 5. This leads me to one of the potential pitfalls of working in (p_M, p_W)-space, namely the perverse effects of overlooking the fact that the health states of the individuals become increasingly correlated as time passes. Such an oversight can lead to very misleading results. The right-hand side panel of Figure 8 shows the evolution of beliefs in Case A for different initial conditions, when no symptoms are observed, but where posteriors are in each period calculated *as if* the health states are independent (rather than correlated). As is clear from the figure, in such a scenario, beliefs continue to increase and always converge to the point $(1, 1)$ in (p_M, p_W)-space. The reason for this result is that when correlation is overlooked, there is a hall of mirrors effect in which each individual can be the "source of their own infection". In other words, after every unprotected encounter, each individual is more likely to be infected than before it (despite showing no symptoms), thereby "increasing" the transmission risk of the partner in the next period in the process. Of course, in reality no individual can in this setting be the source of their own infection and so as times passes and no symptoms are observed, the individuals assign an increasingly high probability to the event that neither is infected.[18]

The model in which individuals ignore the correlation of their health states, can be interpreted as a model of boundedly rational behavior, akin to the persuasion bias introduced by DeMarzo et al. (2003). In that paper, individuals treat each informative signal they receive as being independent information and do not consider the possibility that they may be correlated.

[18] The exception to this, is so-called "ping pong infection" between individuals who can recover from infection, only to be infected by a partner who they themselves have previously infected. But this is not possible in the present setting, since infected individuals cannot recover.

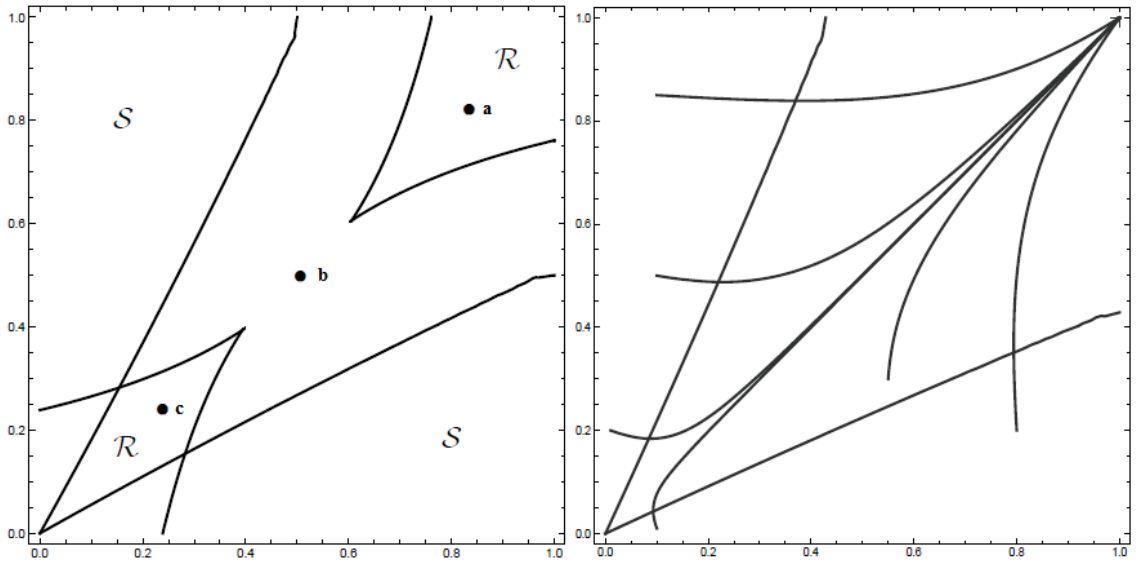

Figure 8: Complications and Pitfalls in (p_M, p_W)-Space.

6. Concurrent Partnerships

The main focus of this paper, is that of a monogamous long-term relationship. This is a natural starting point and also a necessary building block for the analysis of more complex social structures involving multiple interacting individuals. In practice, many individuals engage in overlapping or concurrent partnerships, which by their very nature are more complicated to analyze than the monogamous relationship considered here. There is a large epidemiological literature on the effects of concurrency on HIV propagation. While concurrency is very widespread, its effects on the spread of HIV are still hotly debated (see e.g. Epstein and Morris, 2011). Both empirical and simulation studies make a convincing case that concurrent partnerships are a key ingredient in HIV transmission and they therefore deserve special attention.[19] Some authors argue that concurrency may speed up the spread of infection in a network, but these conclusions come with two important caveats. First, most such studies do not feature decision making. Introducing more channels of infection between individuals will trivially increase the actual amount of infection, ceteris paribus, but rational decision making may counteract such effects.[20] Second, the extent of disease transmission when there are overlapping partnerships, may depend on both endogenous link formation and on the number of sex acts that each of the individuals allocate to each of their partners.[21] While I will not do full justice to these interesting issues, I will briefly outline some salient features that an extended analysis will contain.

For simplicity, I will consider only the simplest extension of the model to include one new individual, namely a lover L. I will consider two separate settings, corresponding to the two possible non-trivial networks that can be formed with three individuals. In the first, the woman is assumed to engage in concurrent relations with the man and the lover, forming a line as indicated in the left-hand side panel of Figure 9. In the other case, all three individuals maintain relations, forming a triangle as indicated in the right-hand side panel. This network is not chosen because it represents a typical case of independent interest. Rather, it is chosen

[19] Epstein and Morris (2011) state that *"If a sufficient fraction of the population has long-term ongoing relationships with more than one person, relatively stable connected sexual networks arise, in which each person's risk is determined not only by his (or her) own behaviour, but also by that of all the others in the network"*.

[20] This critique of biological epidemiology is articulated most explicitly by Geoffard and Philipson (1996).

[21] See Moslonka-Lefebvre et al. (2012) for a micro-founded model of how sex acts are allocated across partners.

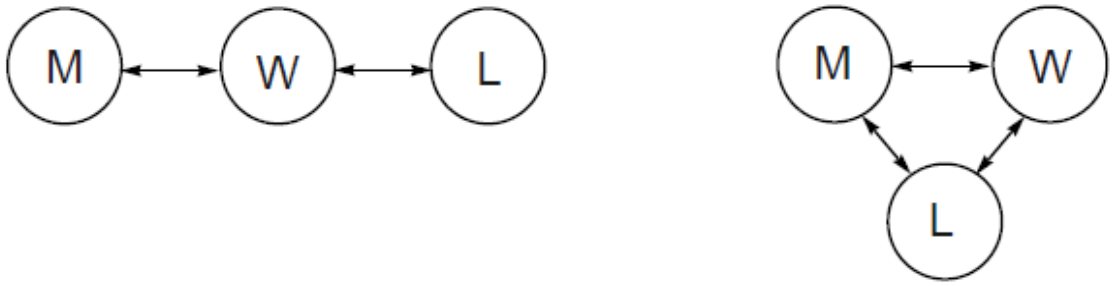

Figure 9: Concurrent Relationships. In left-hand side panel, the woman has concurrent relations with the man and the lover (forming a line); in right-hand side panel, all three have a concurrent relationship (forming a triangle).

because it is the simplest possible network that allows me to make a more general point. In what follows, I will continue to assume that the situation is commonly known by everyone involved.

First, consider the line network. It is important to note that by assumption, there can only be *direct* disease transmission between the man and the woman and between the woman and the lover, respectively. As the man and the lover do not maintain relations, any disease transmission between these individuals must happen *indirectly* and in a subsequent period, via the woman. Also, note that as the individuals are not forward-looking, the only thing that matter for current decision making, are the current transmission risks. In particular, the man will not base his decision on the contemporaneous decisions that the woman and the lover make in their concurrent encounter. The state and message spaces are extended in the obvious manner (to eight possible states). Rather than characterizing the separate evolution of the many state probabilities, I will focus on the transmission risks between the man and the woman, respectively. This will allow me to focus on the effects that adding the lover has on beliefs and transmission risks under unprotected sex between the two concurrent partnerships. Last, note that the situation of the lover is no different in nature than that of the man, vis-à-vis the woman.

In the left-hand side panel of Figure 10, I show some possible paths of the mutual transmission risks r_{MW} and r_{WM} of the man and the woman respectively, under the assumption that no symptoms have been observed in either individual or in the lover. By assumption and to make the comparisons cleaner, the man and the woman initially face identical transmission risks. For comparison, the dashed line in this figure indicates the common transmission risk faced by the couple under monogamy, under the same parameter values and initial conditions.[22]

There are two interesting features worth noting. The first feature is that while the transmission risk that the woman faces from the man is decreasing, as is the case under monogamy, it in fact decreases more sharply, because of the presence of the lover. The reason is simply that as the woman is exposed to additional transmission risk from the lover, she is more likely to become infected and therefore less and less susceptible to disease transmission from the man. Ceteris paribus, this should make the woman more willing to engage in unprotected sex with the man. The second feature, is that the transmission risk faced by the man, is now non-monotone. The reason is that as the woman is now exposed to more transmission risk from her lover, this translates with a lag into a higher risk to the man (through the woman's increased infection probability). Over time, if neither the woman or her lover show symptoms, eventually the perceived transmission risks (for all involved) start decreasing. But at least

[22]Plots of all the state probabilities for each network are available upon request.

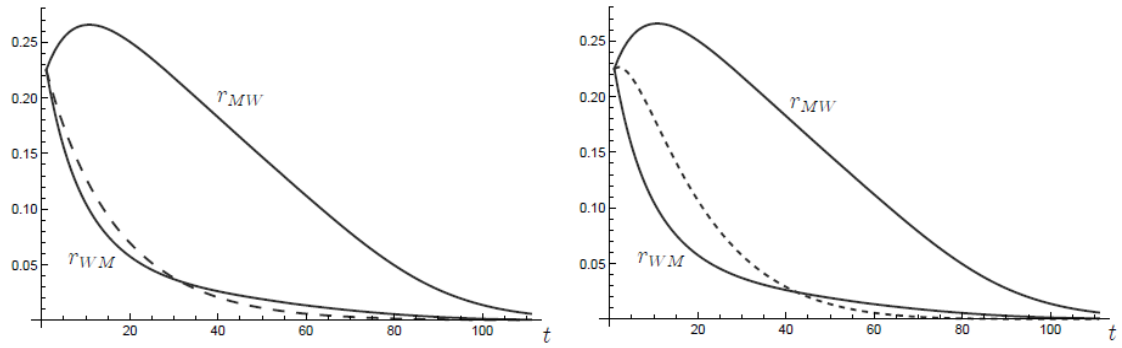

Figure 10: Transmission Risks Under Concurrent Relationships. Left-hand side panel compares case of line with monogamy; right-hand side panel compares case of triangle with line.

initially, ceteris paribus, the presence of the lover may make the man less willing to engage in unprotected sex with the woman. A symmetric argument holds for the lover, because of the woman's unprotected relations with the man. Also note that the total transmission risk that the woman faces is higher than that plotted in Figure 10, because she also faces the additional risk of transmission from her lover.

Next, consider the triangle network. In the right-hand side panel of Figure 10, I again show the paths of transmission risks r_{MW} and r_{WM} from the line network, again under the assumption that no symptoms are observed. Furthermore, the common transmission risks between the man and the woman in the triangle network, are indicated with a dashed line. Note that the transmission risks from the woman to the man and vice versa, are equal as they are now facing perfectly symmetric problems. The addition of the lover to form the triangle network, has two countervailing effects on the inference problems faced by the man and the woman. On the one hand, they now each face higher total transmission risks, compared to the monogamous setting. On the other hand, the addition of the lover, whose health status under unprotected sex becomes increasingly correlated with that of the man and the woman, provides additional information about the health status of all three individuals. In other words, in terms of inference, there is a tradeoff to adding a new individual to the network. The tradeoff is not unlike that facing the individuals under monogamy and the net effect will depend on the parameters (β, σ). Compared to the line network, the woman in the triangle network initially faces higher transmission risks under the triangle network. Again, the transmission risk is non-monotone, although this is less pronounced than under the line network. But interestingly, the transmission risk facing the man is lower in the triangle network, than in the line network. This is because, although the man is now exposed to more risks (from both the woman and the lover), the fact that all three individuals are symptom free and that their health states are correlated, rapidly decreases the perceived transmission risks faced by all three individuals. For other parameter values, the transmission risk in the triangle network could be larger than that of the man in the line network.

It is useful to directly compare the transmission risk from the man to the woman, across networks. In each network, the woman is directly connected to the man and to the lover. The only difference is therefore that in the line network, the woman is the only connection between the two, while in the triangle network, they are also directly connected. As mentioned before, adding this connection between the woman's two partners has two distinct effects. First, the total transmission risk faced by the woman is increased by the addition of the connection, as can be seen in the initial segment of the dashed curve in the right-hand side panel of Figure 10. The reason is that apart from the direct risk posed by being connected to the man and the

lover, the woman now also experiences indirect risks. For example, suppose that the lover is infected but that the man is not. In the line network, the woman can only be infected by the lover (but not by the man). In the triangle network, the lover can directly infect the man, who can in turn infect the woman, even in cases where the lover does not directly infect the woman. Thus the additional connection increases overall infection risks. Second, the addition of the link between the man and the lover means that their health states become more strongly correlated. This means that observing the symptoms of either, becomes more informative about the health state of the other. In other words, adding the connection leads to faster learning. This is seen in the figure by noting that the dashed curve eventually decreases faster than the curve for the woman's transmission risk in the line network.

It is immediately clear that many of these features will extend to more elaborate networks than the ones considered here. The general lesson is that adding individuals to a network has two effects. On the one hand, they add new possibilities of disease transmission. On the other hand, their presence also allows individuals to learn faster, because the health states of connected individuals become increasingly correlated. The net effect will depend not only on the parameters (β, σ) and the initial conditions, but will also be a function of the network topology.[23]

From these considerations, it is clear that the possible equilibrium dynamics become considerably more complicated once additional individuals are added to the network. The best responses in Theorem 1 remain valid in such settings, but there can now be multiple switches between protected and unprotected sex. For example, it is possible to start in a situation with relatively modest transmission risks, which justify all engaging in unprotected sex. Because of the potential for disease transmission, the risks go up, prompting a switch to protected sex. Over time, if no symptoms are observed, there can be an additional switch back to unprotected sex. In fact, such non-monotonicities are not limited to situation in which no symptoms are observed. Suppose that no symptoms have been observed and that all have unprotected sex. Next, suppose that the lover starts showing symptoms, which prompts the man to revise upwards his assessment of being at risk from unprotected sex with the woman and the woman's assessment of being at risk from the lover. In this case, they may all switch to protected sex. Over time, if no further symptoms are observed, the transmission risks will become sufficiently low to justify a switch back to unprotected sex between the man and the woman (but not between the woman and her lover).

Briefly turning to issues of welfare, it is clear how adding concurrent partnerships to the model can create inefficiencies in sexual decision making, in a world with forward-looking behavior. This follows since one individual's decision to engage in unprotected sex today, may increase the transmission risk faced by an indirect partner tomorrow. Unless this effect on third parties is internalized, there is scope for equilibrium over-exposure from the perspective of the utilitarian social planner. The extent of this should be balanced against the inefficiencies created by the least willing individual being pivotal (which will be explored in detail below) and will also depend on how much the individuals discount the future.

As argued by Anglewicz and Kohler (2009), there is evidence to suggest that individuals tend to overestimate their probability of being infected and as a consequence, to underestimate their susceptibility to infection. The persuasion bias studied by DeMarzo et al. (2003) is a possible explanation for this finding. The more connected individuals become in a sexual network, the more their health states become correlated as a consequence of unprotected sex. If an individual has multiple sexual partners who are mutually connected, either directly or indirectly, then something equivalent to persuasion bias may lead individuals to over- or under estimate the transmission risks they face.

[23] For more on disease transmission on sexual networks, see Doherty et al. (2005).

7. WELFARE, ALTRUISM AND SOCIAL CONVENTIONS

In this section, I analyze the social optimality of the equilibrium decisions made by the two individuals. I also consider extensions to situations in which individuals have either altruistic preferences or more extreme self-abnegating preferences. Last, I discuss the role of social conventions in determining the type of sex the couple engages in.

7.1. Socially Optimal Decisions. From the perspective of a utilitarian social planner, there are three distinct sources of inefficiencies in sexual decision making in the context of the present analysis. First, there is the fact that individuals are assumed to be passive learners and thus do not act to forego short term gains in order to make better informed future decisions. In what follows, the social planner will be assumed to adopt the same kind of policies, but with a view to maximize social welfare period by period. Second, decentralized protection decisions typically lead to negative externalities, because individuals do not properly internalize the effects that their decisions have on other individuals. This effect is absent as long as the decision makers are assumed not to be forward looking.[24] The third source of inefficiency, further characterized below, is different in nature. It stems from the assumption that there is no coercion. With this assumption in place, it is the individual who is least inclined to engage in unprotected sex, that is pivotal to the couple's decision. But that means that unprotected sex takes place only when the person who values it the least agrees to it, thus driving a wedge between the equilibrium outcomes and those preferred by a utilitarian social planner. I will now characterize this inefficiency in detail.

Consider a strictly utilitarian social planner, who in each period seeks to maximize the sum of the individuals' welfare levels. Such a planner's objective function is given by

$$W(r_{MW}, r_{WM}) \equiv U_M(r_{MW}) + U_W(r_{WM}) = 2 - \beta(r_{MW} + r_{WM})\pi \tag{40}$$

$$= 2 - \beta\kappa\pi \tag{41}$$

where κ is the transmissive contact rate defined in (1).[25] To ensure that there is a positive social value of protection, I impose the following strengthening of Assumption 1:

Assumption 2. $\beta\pi > 2$.

This assumption means that the net benefits of unprotected sex are not so large, that the planner would wish a susceptible individual to have unprotected sex with a partner who is known to be infected.

While it is trivially true that the planner may wish to make different decisions than the two individuals would on their own, the direction of this discrepancy is interesting. As can be seen from Figure 11, which illustrates the private indifference curves I_M and I_W together with the iso-social welfare curves I_M^* and I_W^*, there are beliefs for which the planner would prefer unprotected sex while one of the individuals would prefer protected sex. That is, the social planner would prefer *more* unprotected sex than what would result from decision making by the couple (when each individual seeks only to maximize his or her own welfare). The reason for this result can be seen by considering a set of beliefs (p_M, p_W) that is just within the set $\mathcal{S}_i$, $i = M, W$, but far from the set $\mathcal{S}_j$. In this case, the couple will engage in protected sex because i can insist on it, but i's welfare loss from engaging in unprotected sex is very small compared to the welfare gain for j from unprotected sex. The social planner would in this

[24] To properly model this possibility, one would have to assume forward-looking behavior on the part of both the individuals and the planner. This is because under passive learning, there are no contemporaneous externalities in the present model. This issue is discussed further in Toxvaerd (2014).

[25] Note that in the special case $p_M = p_W = I$, $\kappa = 2I(1-I) = 2Var(H)$, where H is the health status of a randomly chosen individual.

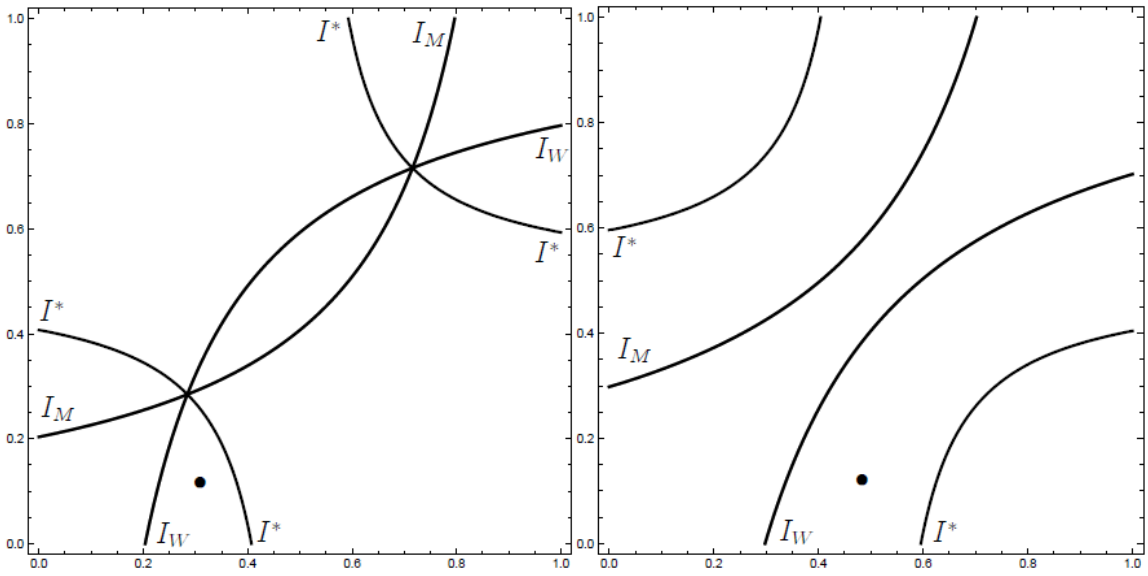

Figure 11: Socially Optimal Versus Equilibrium Decisions in (p_M, p_W)-Space.

situation prefer that the couple engage in unprotected sex, despite i's wishes. Two such points are indicated on the figure.

To formally analyze the differences between equilibrium outcomes and socially optimal outcomes, it is instructive to determine the conditions under which the planner and the individuals decide on unprotected sex, respectively. The planner will prefer unprotected sex when

$$W(r_{MW}, r_{WM}) > 0 \tag{42}$$

which can be rewritten as the condition

$$\frac{1}{2}(r_{MW} + r_{WM}) \leq \bar{p} \tag{43}$$

Turning to the couple, they will engage in unprotected sex when

$$\min\{U_M(r_{MW}), U_W(r_{WM})\} > 0 \tag{44}$$

which can be rewritten as the condition

$$\max\{r_{MW}, r_{WM}\} \leq \bar{p} \tag{45}$$

From these inequalities and the fact that $\max\{r_{MW}, r_{WM}\} \geq \frac{1}{2}(r_{MW} + r_{WM})$, it follows that the individuals' condition for engaging in unprotected sex is more difficult to satisfy than that of the planner. Furthermore, $\max\{r_{MW}, r_{WM}\} = \frac{1}{2}(r_{MW} + r_{WM})$ for $r_{MW} = r_{WM}$. This means that on the diagonal in (p_M, p_W)-space, equilibrium decisions are always socially optimal. In particular, this means that any of the states $(1, 0, 0, 0)$ and $(0, 0, 0, 1)$ are valued as much by the planner as they are by the individuals. The reason for this finding is simply that for $r_{MW} = r_{WM}$, the two individuals face identical problems and thus their equilibrium decisions coincide. For this reason, their decisions also coincide with those preferred by the planner.

Define the safe set under utilitarian social planning as

$$\mathcal{S}^* \equiv \{(r_{MW}, r_{WM}) \in [0, 1]^2 : W(r_{MW}, r_{WM}) \leq 0\} \tag{46}$$

I can now state the following result:

Theorem 3. *The safe sets are ordered so* $\mathcal{S}^* \subseteq \mathcal{S}$.

Because there is no coercion, there can never be "too much" exposure to infection in equilibrium. This is so because if either individual prefers protected sex, then no unprotected sex will take place. On the other hand, the social planner, who takes into account the welfare of both individuals, may well prefer that unprotected sex takes place, since the increase in welfare of the willing partner outweighs the loss in welfare of the unwilling partner. Since the safe sets are different under the two scenarios, for a given initial belief, the equilibrium path may end in a different steady state than would have been the case under socially optimal decisions. Last, it is worthwhile noting that over time, as correlation between the two individuals' health states increases, the inefficiencies in equilibrium decisions tend to disappear.

It should be emphasized that the above results are very specific to a social planner that is utilitarian, i.e. one that seeks to maximize the sum of the individuals' utilities. If one instead considers a Rawlsian social planner, i.e. one that seeks to maximize the utility of the worst-off individual, then the joint decisions under selfish preferences are in fact socially optimal, as they favor the individual who is most susceptible at the expense of the one that is least susceptible.

Welfare and Equilibrium Decisions with Transfers. Without transfers, there is no way for the individuals to align their preferences and achieve the socially efficient decision. For that reason, I now relax the assumption that the individuals cannot make in transfers. As noted by Philipson and Posner (1993) and further discussed by Greenwood et al. (2013), such transfers should be interpreted in the widest possible way, to include the exchange of both such "commodities" as companionship and material gifts. Luke et al. (2011) collect empirical evidence of gift giving between sexual partners.

If individuals have quasilinear preferences (i.e. preferences that are linear in the transfer), the socially optimal decisions can be obtained without coercion, simply by allowing side payments between the individuals, i.e. by allowing an individual to compensate its partner for engaging in more unprotected sex than desired. A transfer function that will achieve this is e.g. a transfer from individual i to individual j given by

$$\tau(r_{ij}, r_{ji}) = U_i(r_{ij}, r_{ji}) \tag{47}$$

This transfer function keeps individual i on his or her reservation utility, since

$$U_i(r_{ij}, r_{ji}) - \tau(r_{ij}, r_{ji}) = 0 \tag{48}$$

while individual j now acts with a view to maximize

$$U_j(r_{ji}, r_{ij}) + \tau(r_{ij}, r_{ji}) = W(r_{ij}, r_{ji}) \tag{49}$$

In other words, social welfare is maximized once transfers are allowed.

For concreteness, consider the following numerical example, in which the health states are assumed independent. Let the infectiousness be $\beta = 0.6$ and let the benefit-cost ratio be $\pi = 5$. With $p_M = 0.2$ and $p_W = 0.5$, it is easy to verify that $U_M(r_{MW}, r_{WM}) = -0.2$, $U_W(r_{WM}, r_{MW}) = 0.7$ and $W(r_{MW}, r_{WM}) = 0.5$. In other words, with these parameter choices and for these beliefs, the man would prefer protected sex, while the woman would prefer unprotected sex. Without coercion, the man's choice determines the type of sex that the couple engages in and thus they choose protected sex. But the planner clearly values unprotected sex over protected sex, since the net social welfare resulting from unprotected sex is positive. The socially optimal outcome can be achieved by implementing any transfer $\tau(r_{MW}, r_{WM}) \in$

$[0.2, 0.7]$ from the woman to the man. For any such transfer, each individual's participation constraint is satisfied and the socially optimal outcome is achieved.

Although the presence of transfers may seem to achieve the socially optimal outcome, this finding comes with an important caveat, for it is based on a strictly utilitarian welfare measure. While utilitarianism is pervasive in economic analysis, many people are uncomfortable with using it in contexts of health. For in its strictest interpretation, utilitarianism sanctions the harvesting of organs from healthy individuals, even without their consent. It can be argued that the social planner should instead pursue a Rawlsian objective, which is exactly what is achieved when there is no coercion and no transfers. In conclusion, whether transfers should be allowed, reduces to a question of society's view on how to aggregate individuals' welfare. In Germany, recently proposed legislation seeks to ban prostitutes from offering unprotected sex to their costumers, essentially stipulating that no transfer will make this service available.[26]

7.2. Altruistic Preferences. So far, I have maintained the assumption that the individuals are strict utility maximizers, with no regard for their partner's well-being. This may be a strong assumption and it is worthwhile discussing what would happen if one relaxes it. Altruism in the context of sexual decision making has been discussed by Philipson and Posner (1993) and Gong (2015). O'Dell et al. (2008) argue that *"In the context of sexual risk behavior, both self-interest and concern for others may play a role in the decision to practice safer sex"*. As will become clear in what follows, while it is certainly true that altruism may change sexual decision making, the direction of the change may be somewhat surprising. In the large body of literature on sexually transmitted diseases, it is quite common to assert that particular outcomes are the effects of altruism, without explicitly verifying that altruistic motives would lead to the observed outcomes. For example, Duffin (2004) reports that in HIV education, *"[Altruism] describes the responsibility of a person with HIV to avoid transmitting the virus"*.[27] In contrast, I will instead assume that individuals have altruistic preferences and then see where this assumption leads. As will become clear in what follows, altruism can have surprising effects on the decision to engage in voluntary transmissive sexual behavior.

To model altruism, suppose that the objective of an individual $i = M, W$ is to maximize

$$V_i(r_{ij}, r_{ji}) \equiv U_i(r_{ij}) + AU_j(r_{ji}) \tag{50}$$

where $A \in [0, 1]$ is the degree to which individual j's well-being is take into account by individual i. This formulation, which is the standard non-paternalistic approach used in the literature, nests the two extremes of purely selfish preferences on one hand and the social preferences of the utilitarian social planner on the other.

The best responses of the two individuals and their aggregation to joint decisions in Theorem 1 are modified in a straightforward way. With these preferences, the couple will engage in unprotected sex when

$$\min\{V_M(r_{MW}, r_{WM}), V_W(r_{WM}, r_{MW})\} > 0 \tag{51}$$

which can be rewritten as the condition

$$\max\left\{\frac{r_{MW} + Ar_{WM}}{1 + A}, \frac{r_{WM} + Ar_{MW}}{1 + A}\right\} \leq \bar{p} \tag{52}$$

It is straightforward to verify that condition (52) coincides with condition (43) for $A = 0$ and

[26] Condoms Compulsory in German Sex Industry Under New Law, The Guardian, February 3rd, 2016.

[27] Similarly, Kerwin (2014) states that *"[...] purely self-interested people should see little or no marginal cost from further risky sex if they are already infected, while altruistic people would want to take measures to protect their prospective partners. A parallel logic applies to those who learn they are HIV-negative"*.

with condition (45) for $A = 1$.
Define the safe set under altruistic preferences as

$$\mathcal{S}^A \equiv \{(r_{MW}, r_{WM}) \in [0,1]^2 : \min\{V_M(r_{MW}, r_{WM}), V_W(r_{WM}, r_{MW})\} \leq 0\} \tag{53}$$

As can be seen from equation (52), the safe and risky sets under altruism are less straightforward than those under selfish decision making and has similarities with those under socially optimal decision making. A direct comparison of the conditions (43), (45) and (52) shows that selfish individuals are least inclined towards unprotected sex, followed by the altruistic individuals, who are in turn less inclined towards unprotected sex than what is socially optimal from a utilitarian perspective. Formally, if $r_{MW} \geq r_{WN}$, then

$$r_{MW} \geq \frac{r_{MW} + Ar_{WM}}{1+A} \geq \frac{r_{MW} + r_{WM}}{2} \geq \frac{r_{WM} + Ar_{MW}}{1+A} \geq r_{WM} \tag{54}$$

I thus have the following result:

Theorem 4. *The safe sets are ordered so* $\mathcal{S}^* \subseteq \mathcal{S}^A \subseteq \mathcal{S}$.

At first glance, it may seem surprising that in this model, altruistic behavior implies that there is more, rather than less, unprotected behavior and hence a higher probability of disease transmission. The key to understanding this result, is to recall that unprotected sex must be voluntary. This means that even if an individual who is very likely to be infected has the well-being of his highly susceptible partner in mind when considering his options, it is in fact the decision of the highly susceptible individual that counts. And as long as the low risk individual takes the well-being of the high risk individual into account, she will be more inclined towards unprotected sex than a purely selfish individual. The iso-welfare curves and the indifference curves under altruism are depicted in Figure 12. At the indicated points, selfish individuals would choose to engage in protected sex, while altruistic individuals would opt for unprotected sex. It should be noted that for $r_{MW} = r_{WM}$, altruistic decisions coincide with those under selfish preferences and are socially optimal.

As noted above, the finding that the planner would prefer there to be more unprotected sexual relations than in equilibrium, depends on the restriction to passive learning. More generally, in a setting with concurrent partnerships and active learning (i.e. with forward-looking behavior), unprotected sex will have negative external effects on third parties. This implies that a social planner would prefer less unprotected sex than in the present benchmark setting. This raises an interesting possibility in the context of altruism. There is strong evidence, both from the laboratory and from field experiments, that individuals display more altruism towards *specific* (or identified) individuals, than towards *abstract* (or unidentified) individuals (see Small and Loewenstein, 2003 and references therein). This can lead to counter-intuitive consequences. Suppose that while individuals are altruistic, the altruism is only directed towards their current partner(s) and not towards the general population. To make the point as clean as possible, suppose that $A = 1$. In this case, altruism may in fact *decrease* social welfare, because each individual, in trying to take into account the well-being of only their direct partner(s), may agree to too much unprotected sex, from the perspective of the social planner (who factors in the future negative externalities imposed on third parties). Such directed altruism seems much more realistic than a more generalized form of altruism. One therefore faces the surprising prospect that an increase in such (directed) altruism can make society as a whole worse off!

7.3. Social Conventions. In this paper, I have adopted the social convention that unprotected sex must be consensual. In particular, this means that the least willing individual is

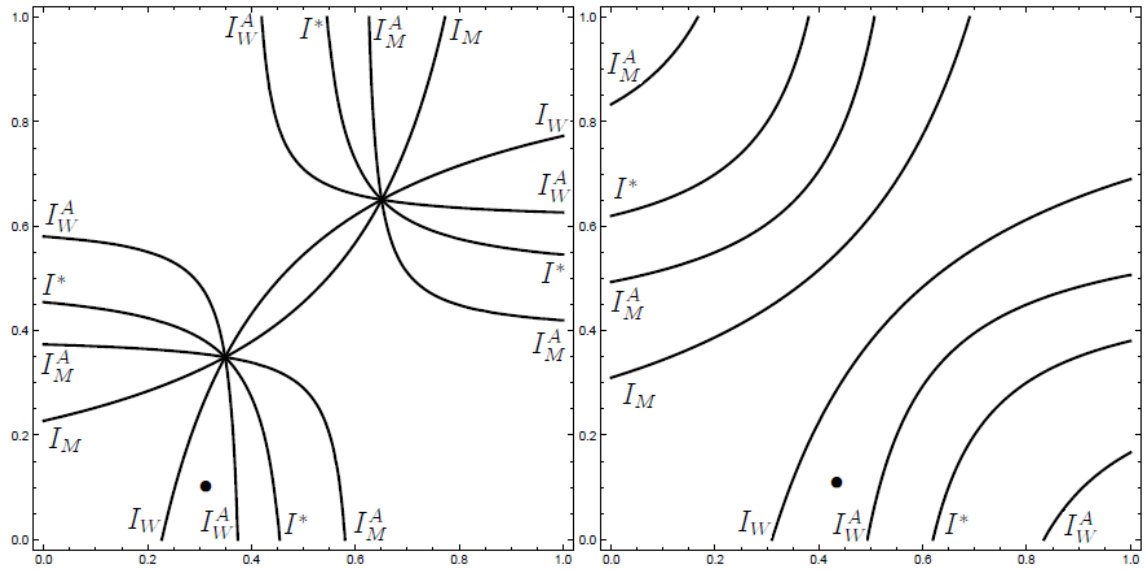

Figure 12: Indifference Curves and Iso-Welfare Curves with Altruism in (p_M, p_W)-Space.

pivotal to the joint decision to engage in unprotected sex. This modeling choice is of course arbitrary to some extent, reflecting a personal value judgment as much as it reflects reality. For that reason, I will now discuss the alternative assumption where

$$a_t(r_{ij}, r_{ji}) = \max\{a_t^i(r_{ij}), a_t^j(r_{ji})\} \tag{55}$$

This assumption describes situations in which there is coercion. While this assumption may seem extreme, it serves as a useful benchmark for comparison.[28]

It is immediately clear that the dynamic aspects of the analysis are qualitatively unchanged, with the least susceptible individual's threshold replacing that of the most susceptible individual. This is seen clearly from the last two inequalities in (54). Turning to the issue of welfare and altruism, all the central inequalities are reversed. This means in particular that from a social perspective, there is in this case too much unprotected sex. This is because the least willing individual will now be forced to engage in unprotected sex against his will, in order to please his partner. In this case, altruism has the effects that one may intuitively have expected to hold in general, namely that the extent of unprotected sex is reduced. The reason is simply that when the least susceptible individual is pivotal, then altruism will induce this individual to factor in the well-being of the more susceptible (and therefore less willing) individual. This leads to less transmissive joint decisions. In the previous section, I noted that when transfers are allowed, the equilibrium decisions with side payments achieve the social optimum. This is still the case under coercion, but here payments take a more sinister form. In particular, payments now take the form of an inducement from the least willing partner to the more willing partner, in order to avoid having to engage in unprotected sex.

To more clearly draw out the consequences of the different preferences and social conventions, note that under both selfish, altruistic and socially optimal behavior, joint decisions are always (weakly) increasing in the transmission risks r_{ij} and r_{ji}. This means that I can rank the different scenarios in terms of how quickly the individuals switch from protected to unprotected behavior. Making use of (54), Figure 13 illustrates the different conditions, in the case in which

[28] In 1 Corinthians 7:5 of the King James Bible, the married couple is admonished thus: *"Defraud ye not one the other, except it be with consent"*. This is generally interpreted to mean that the couple should only *abstain* when doing so is consensual.

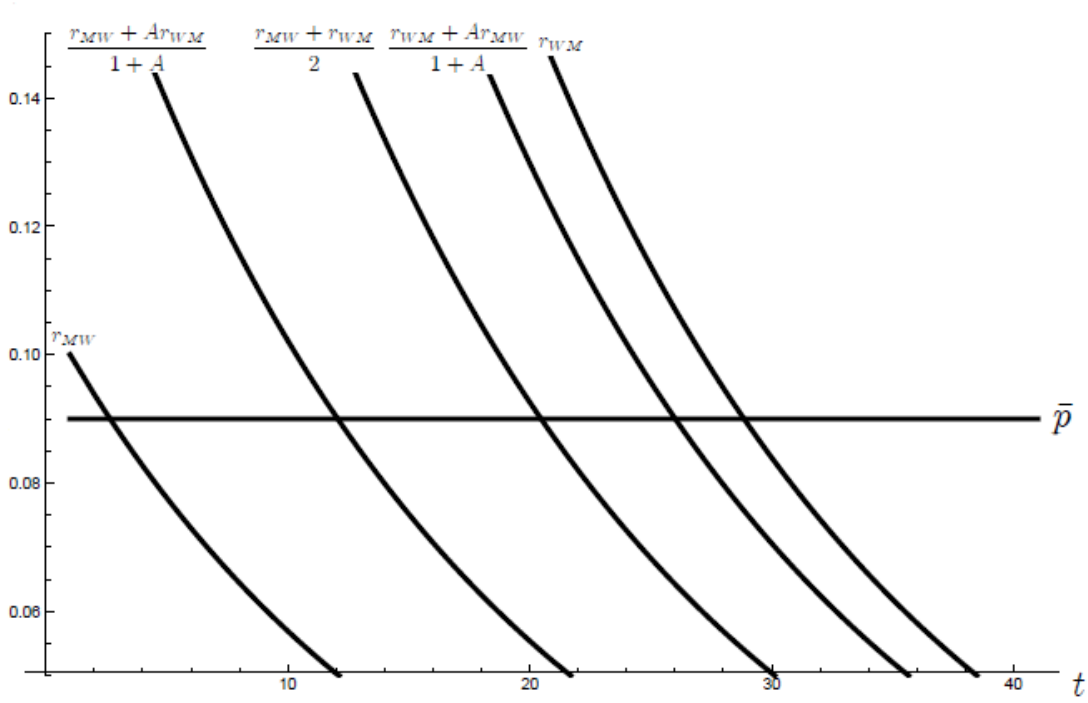

Figure 13: Switching Conditions When No Aymptoms Are Observed. Curves from left to right represent (i) selfish preferences under coercion, (ii) altruistic preferences under coercion, (iii) the social optimum, (iv) altruistic preferences without coercion and (v) selfish preferences without coercion.

no symptoms are observed. The middle curve shows the relevant condition for the socially optimal switch from unprotected to protected sex. The outermost curves show the conditions under coercion (left) and no coercion (right) respectively, each under selfish decision making. In between those curves and that of the socially optimal case, are the relevant curves under coercion (left) and no coercion (right) respectively, but under altruistic decision making. As the altruism parameter A tends to one, these latter curves approach (and eventually coincide with) the middle curve. Conversely, as A tends to zero, the curves approach (and eventually coincide with) the outermost curves. It is clear from this figure that as altruism increases, equilibrium decisions necessarily approach the socially optimal decisions. Having said that, the increase in social welfare comes about in different ways, depending on the social norm. Under coercion, altruism increases social welfare by delaying the switch from protected to unprotected sex (thereby decreasing overall transmission). But under no coercion, altruism increases social welfare by switching to unprotected sex sooner, thereby increasing overall transmission.

I have for clarity only considered the two extremes in which the individual either least or most susceptible to infection, is pivotal for the decision to engage in unprotected sex. In general, the outcome will depend on social (i.e. cultural) norms, idiosyncratic power relations within the couple and the bargaining process that aggregates preferences into decisions. But it should be clear that the more one or the other part is favored, the more one moves towards the extremes I have described.

It is of course an empirical questions whether one norm or the other is prevalent in a given society (or indeed, within the context of a particular relationship). Some work has been carried out on this topic by, Muldoon et al. (2014). They find that there is wide variability across couples within the same society, with some couples in their sample following joint decision making and other couples following decision making by one individual only. They show that

when couples make joint decisions, they are significantly more likely to use condoms. They also emphasize the role of power relations within the couple in determining the type of decision making and survey a large literature on the topic. These findings suggest that policies designed to shift cultural norms of sexual decision making may be a fruitful avenue for influencing aggregate patterns of exposure and disease transmission.

7.4. Self-Abnegating Preferences. As an alternative to the type of altruistic preferences treated above, several authors consider more extreme forms of altruism, where individuals self-sacrifice their own utility in order to increase the well-being of their partner (see e.g. Chen and Kempe, 2008 and references therein). In particular, individuals are assumed to maximize an objective of the following type:

$$\tilde{V}_i(r_{ij}, r_{ji}) \equiv (1-A)U_i(r_{ij}) + AU_j(r_{ji}) \tag{56}$$

Here, $A \in [0,1]$ denotes the weight given by an individual to the partner's well-being and simultaneously, the degree to which the individual is self-abnegating. With such preferences, the couple engage in voluntary unprotected sex when

$$\min\{(1-A)U_M(r_{MW}) + AU_W(r_{WM}), (1-A)U_W(r_{WM}) + AU_M(r_{MW})\} > 0 \tag{57}$$

This condition can be rewritten as

$$\min\{U_M(r_{MW}) - A\left[U_M(r_{MW}) - U_W(r_{WM})\right], U_W(r_{WM}) + A\left[U_M(r_{MW}) - U_W(r_{WM})\right]\} > 0 \tag{58}$$

In this case, it is clear that as the altruism parameter A increases, the least susceptible individual become increasingly inequality averse, while the most susceptible individual becomes increasingly inequality loving. This can be seen by recalling that if $r_{MW} \leq r_{WM}$, then $U_M(r_{MW}) \geq U_W(r_{WM})$. In the extreme case where $A = 1$, the two individuals essentially swap preferences. As the individuals become completely self-abnegating, the least susceptible of the couple would block unprotected sex, exactly when the most susceptible individual would have done so under completely selfish preferences.
In terms of transmission risks, the condition for voluntary unprotected sex becomes

$$\max\{r_{MW} + A\left[r_{WM} - r_{MW}\right], r_{WM} - A\left[r_{WM} - r_{MW}\right]\} \leq \bar{p} \tag{59}$$

To make things interesting, I impose the following:

Assumption 3. $\max\{r_{MW}, r_{WM}\} \geq \bar{p}$.

Assumption 4. $\frac{1}{2}(r_{MW} + r_{WM}) \leq \bar{p}$.

These assumptions ensure that there are cases in which purely selfish individuals, would make different choices than self-abnegating individuals.

The joint decisions under this kind of preferences, when unprotected sex must be voluntary, are illustrated in Figure 14. The V-shaped segment in the figure drawn in heavy line, represents the condition for the joint decision to engage in unprotected sex.

As can be seen from the figure, for $A = 0$, the individuals are entirely selfish and the man would block unprotected sex, as $r_{MW} > \bar{p}$. As the man puts increasingly high weight on the utility of his partner, he become increasingly willing to engage in unprotected sex. Note that for the special case $A = 1/2$, the individuals in fact care as much about themselves as they do about their partner and hence their joint decision coincides with that chosen by the utilitarian social planner. But note that when A increases beyond $1/2$, the couple becomes less and less willing to engage in unprotected sex. The reason is that in this case, the woman puts more

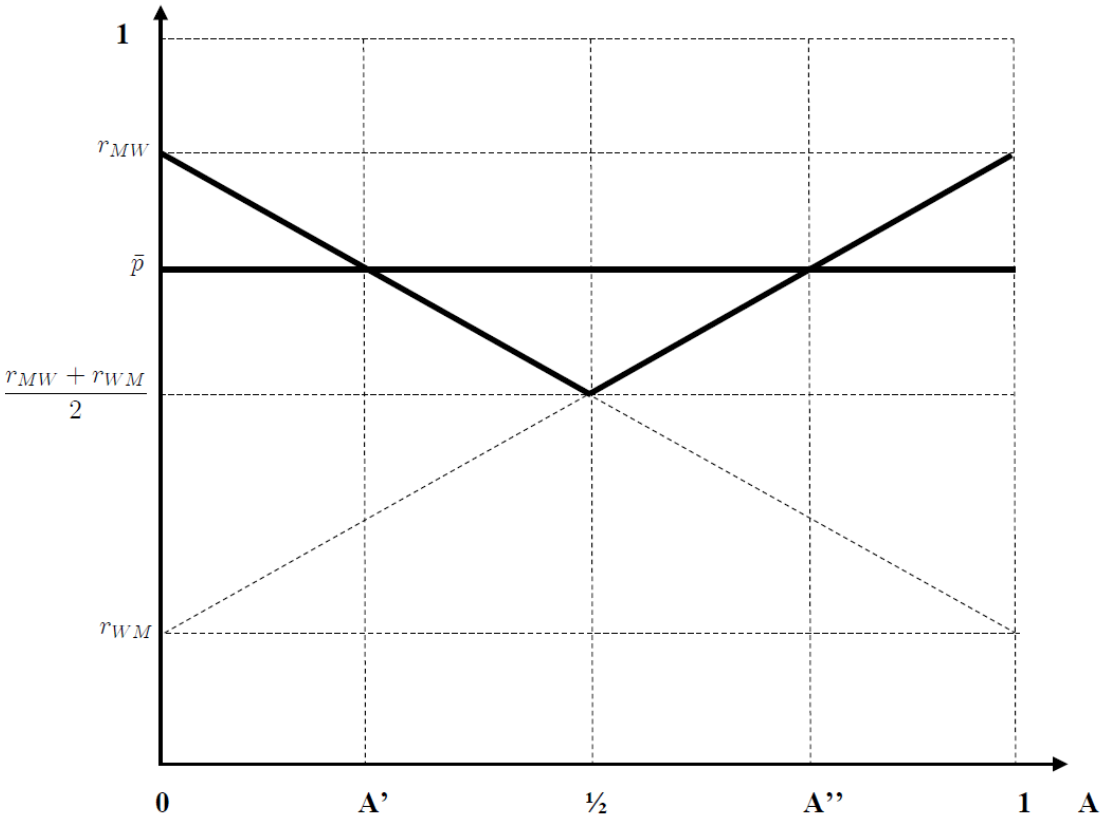

Figure 14: Joint Decisions Under Self-Abnegating Preferences.

weight on the man's utility than on her own and therefore starts acting more and more as would a man under selfish preferences. She therefore becomes increasingly unwilling to engage in unprotected sex. In the extreme case where $A = 1$, the two individuals have essentially swapped preferences and therefore block unprotected sex exactly as would individuals under selfish preferences.

Last, note that in the extremes, i.e. for completely selfish and for completely self-abnegating preferences, the joint decision differs from that chosen by the utilitarian social planner, but coincides with that chosen by the Rawlsian social planner. In contrast, in the middle area where $A \in [A', A'']$, the individuals maximize the sum of their (material) utilities. It is interesting to note that for extremely self-abnegating preferences, the joint decision will coincide with those preferred by either selfish individuals or by a Rawlsian social planner. In contrast, extremely altruistic (but not self-abnegating) preferences would lead to joint decisions that coincide with the decision chosen by a utilitarian social planner.

For completeness, it should be mentioned that with self-abnegating preferences and coercion, the most susceptible individual may act to ensure that the couple engage in unprotected sex, in order to please the less susceptible partner.

8. Discussion

In this paper, I presented and analyzed a simple model of sexual decision making when individuals are faced with an infectious but possibly asymptomatic disease. The disease can spread via unprotected sexual behavior. I showed that behavior and learning interact in interesting ways. In particular, because individuals learn about their infection state, which evolves endogenously over time as a function of the decisions on protected versus unprotected sex, individuals may rationally assign either increasing or decreasing probabilities to being infected, even without

observing symptoms. The non-monotonicity is brought about by the interaction of exposure from unprotected sex and a lack of symptoms. I furthermore extend the analysis to some simple settings with concurrent partnerships. Next, the analysis showed that when unprotected sex happens only on a voluntary basis, there is too little such sex from the perspective of a utilitarian social welfare criterion. This follows from the observation that under voluntary sex, the most at-risk individual can veto the joint decision of the couple with a view to maximizing own welfare only. I showed that when individuals become increasingly altruistic, the extent of voluntary unprotected sex increases. Last, I considered the problem of identifying the index case as a function of the timing of symptoms.

While the model presented in this paper has yielded a number of results, there is scope for extending the analysis in several different directions. First, it is interesting to go further in the analysis of concurrent partnerships. In practical terms, this will involve embedding each decision maker in a structured population model such as a network. Ideally, such networks would have both stable components (to capture long term relationships) and random formation of additional connections (to capture the possibility short term or opportunistic sexual relations). In such a setting, the correlations between the individuals' health states, their decisions and the resulting disease dynamics will depend delicately on the network structure of the population.

Second, it is interesting to relax the assumption that individuals are passive learners and to determine the dynamics of the model, when individuals engage in optimal (or strategic) experimentation. It is not clear whether optimal (or equilibrium) experimentation would lead to less or to more protected sex, as there is a tradeoff to be considered. On the one hand, protected sex makes symptoms (or the absence of such) a more precise signal of past health states. This would suggest that an experimenter would opt for less unprotected sex in order to learn. On the other hand, under unprotected sex, the health states of the two individuals are more likely to be the same, and thus one may learn something about one's own health state by observing the symptom history of one's partner. It seems that resolving this tradeoff would rely in a delicate manner on the two central parameters (β, σ). From the present analysis it appears that the value of experimentation is decreasing in the correlation of the individuals' health states and thus would be most relevant at early stages of the relationship.

Third, I have assumed that once an individual is infected, he or she cannot recover. This assumption reflects the fact that HIV/AIDS is the main motivating example for the present model. To adapt the analysis to other diseases in the SIS class, one can extend the model in a straightforward manner by introducing the possibility that an infected individual will recover in each period.

Fourth, I have assumed throughout that while an infected individual could appear healthy, symptoms are a perfect signal that the individual is in fact infected. In other words, I have allowed for false negatives but not for false positives. Allowing for false positive changes the dynamics of the model in interesting ways. Some discussion of such a setting is offered in Online Appendix C.

Fifth, in the present analysis, I disregard all issues of private information, by maintaining the assumption that the infection probabilities of the individuals are common knowledge. In practice, each individual will have superior information about their sexual histories (and indeed, about their relations with third parties). Under such a setting, learning would be immensely more complicated, because there would be private updating (i.e. for the same symptom observations, the individuals would update differently). Furthermore, in such a setting, the individuals would have an incentive to try to elicit their partner's private information by observing their willingness to engage in unprotected sex. Similarly, they would have an incentive to influence their partner's beliefs to their own advantage.

Last, the analysis has assumed that the individuals cannot (or are not willing to) take tests to reduce uncertainty. In many countries, the majority of people do not have access to reliable

testing (see e.g. Anglewicz and Kohler, 2009) or if they do, may not have confidence in their accuracy (see e.g. De Paula et al., 2013).[29] Testing has been the object of both theoretical and empirical analysis in different contexts by Philipson and Posner (1993), Boozer and Philipson (2000), Mechoulan (2004), Sen (2004), Wilson (2010), Gersovitz (2011), Gong (2015) and Ostermann et al. (2015). In these papers, individuals test to either improve the match with other partners, or to commence the appropriate treatment in case they are found to be infected. The theories brought forward by these studies are not dynamic in nature and could in principle be applied in each period of the present analysis. While these papers give convincing arguments for the value of testing, the fact remains that HIV has spread in populations regardless of the widespread availability of inexpensive HIV tests. Furthermore, most people have only a few HIV tests over their lifetime. This suggests that the incentive to test and the tradeoffs involved in testing are not yet fully understood. An additional consideration is that the testing literature focuses on the benefits of treatment, mostly from the perspective of the value of information. But there are a number of generally unappreciated costs involved in testing, which may act to deter individuals from taking tests. Such costs include (i) the inability of an individual to obtain health insurance when a positive HIV status is known and must be disclosed, (ii) the legal sanctions imposed on individuals who know they are HIV positive but engage in unprotected sex and put other's health at risk, (iii), disclosure requirements towards new partners, underpinned by the treat of criminal prosecution (see O'Byrne et al., 2013), (iv) the psychological discomfort involved in knowing with certainty that one is infected (rather than merely suspecting it) and (v) the signal that taking an HIV test may send to a long-term partner (i.e. an admission of possible infidelity); similarly, asking one's partner for a test may be construed as an accusation of infidelity. These and other reasons are described in Johnston (1995) and Fortenberry et al. (2002). Once such costs are considered, potential testers face a real tradeoff and may thus opt to remain uncertain.

[29] De Paula et al. (2013) report that in a study from Malawi, a large majority of men that tested positive for HIV but were asymptomatic, assigned a probability less than one to being infected two years later. Surprisingly, a plurality of these assigned probability zero to being infected.

A. Beliefs Over States

Let the prior beliefs be given by $\{p_1, p_2, p_3, p_4\}$. The posterior beliefs under unprotected sex are as follows:

$$P(x_t = s_1|y_t = r_1) = \theta p_1 \tag{60}$$
$$P(x_t = s_2|y_t = r_1) = \theta(1-\beta a_{t-1})(1-\sigma)p_2 \tag{61}$$
$$P(x_t = s_3|y_t = r_1) = \theta(1-\beta a_{t-1})(1-\sigma)p_3 \tag{62}$$
$$P(x_t = s_4|y_t = r_1) = \theta(1-\sigma)^2[p_4 + \beta a_{t-1}p_2 + \beta a_{t-1}p_3] \tag{63}$$
$$P(x_t = s_1|y_t = r_2) = 0 \tag{64}$$
$$P(x_t = s_2|y_t = r_2) = \frac{p_2(1-\beta a_{t-1})\sigma}{p_2(1-\beta a_{t-1})\sigma + (p_4 + \beta a_{t-1}p_3)\sigma(1-\sigma)} \tag{65}$$
$$P(x_t = s_3|y_t = r_2) = 0 \tag{66}$$
$$P(x_t = s_4|y_t = r_2) = 1 - \frac{p_2(1-\beta a_{t-1})\sigma}{p_2(1-\beta a_{t-1})\sigma + (p_4 + \beta a_{t-1}p_3)\sigma(1-\sigma)} \tag{67}$$
$$P(x_t = s_1|y_t = r_3) = 0 \tag{68}$$
$$P(x_t = s_2|y_t = r_3) = 0 \tag{69}$$
$$P(x_t = s_3|y_t = r_3) = \frac{p_3(1-\beta a_{t-1})\sigma}{p_3(1-\beta a_{t-1})\sigma + (p_4 + \beta a_{t-1}p_2)\sigma(1-\sigma)} \tag{70}$$
$$P(x_t = s_4|y_t = r_3) = 1 - \frac{p_3(1-\beta a_{t-1})\sigma}{p_3(1-\beta a_{t-1})\sigma + (p_4 + \beta a_{t-1}p_2)\sigma(1-\sigma)} \tag{71}$$
$$P(x_t = s_1|y_t = r_4) = 0 \tag{72}$$
$$P(x_t = s_2|y_t = r_4) = 0 \tag{73}$$
$$P(x_t = s_3|y_t = r_4) = 0 \tag{74}$$
$$P(x_t = s_4|y_t = r_4) = 1 \tag{75}$$
$$\theta \equiv \frac{1}{p_1 + p_2(1-\beta a_{t-1})(1-\sigma) + p_3(1-\beta a_{t-1})(1-\sigma) + [p_4 + \beta a_{t-1}p_2 + \beta a_{t-1}p_3](1-\sigma)^2} \tag{76}$$

B. Who is the Index Case? Decoding and the Most Likely Path

In the main body of the paper, the dynamics of beliefs and decisions were analyzed. This analysis relied on filtering, i.e. on determining the distribution over the current (hidden) state, conditional on available information. In the present section, I treat a different but related problem, namely that of determining, given available information, the most like path through time of the state variable. This problem, known as decoding, is useful when one or both individuals in the couple shows symptoms of infection. Suppose that the symptom onset dates are available and one wishes to determine which individual is the most likely source of infection (in the language of epidemiology, the *index case*). Decoding helps answer this question, by determining the probabilities, given the initial conditions and the symptom history, of different sequences of state transitions. This in turns allows me to determine the relative likelihood of histories where one or the other individual is the index case (or indeed the history in which both were infected from the outset). Magder and Brookmeyer (1993) and Teunis et al. (2013) develop statistical methods to determine the likely source of the infection in a seroconcordant infected couple, while Groendyke et al. (2011) treat a similar problem from a theoretical perspective in the context of an SEIR epidemic on a random network.

Given a set of observations $(y_1, y_2, ...)$, the probability that a given path $(x_0, x_1, x_2, ...)$ has materialized can be found as the product of the transition probabilities along the path $(x_0, x_1, x_2, ...)$ and the probabilities that the observed signals where emitted in the states along the path. At the end of this Appendix, the probabilities for all possible paths are calculated for

the case of two observations. While it is possible to find the path probabilities by brute force when the paths are short, it quickly becomes intractable and alternative procedures must be used. A useful approach to calculating these probabilities is known as the Viterbi algorithm, which is given by the following recursion:

$$\max_{\{x_k\}_{k=1}^{t-1}} P(x_0, x_1, ..., x_{t-1}, X_t|y_{1:t}) = \nu P(y_t|X_t) \max_{x_{t-1}} \left\{ P(X_t|x_{t-1}) \max_{\{x_k\}_{k=1}^{t-2}} P(x_1, ..., x_{t-2}, x_{t-1}|y_{1:t-1}) \right\} \tag{77}$$

where the normalizing factor ν is defined in (20).

The Viterbi algorithm is the well-known dynamic programming recursion, applied to the distribution across states.[30,31] The problem of decoding is illustrated in Figure 15, which shows the possible state transitions with two periods. Note that for clarity, only state transitions that have positive probability are indicated. As can be seen from the diagram (a so-called *trellis* diagram), the number of possible paths that the individual has to keep track of grows very fast. For a history with n observations, there are 4^{n+1} different possible paths through the system.

To understand the workings of the Viterbi algorithm, consider a path that ends in state x_t. The most likely path that ends in state x_t can be decomposed into a most likely path to some state x_{t-1} and a subsequent transition from x_{t-1} to x_t. The state x_{t-1} that will become part of the most likely path to x_t, is simply that which maximizes the likelihood of that path.

In practical terms, the most likely path is found by determining for each time t and each state x_t, the *most likely originator*, i.e. the state from which a transition into state x_t is most likely (given the distribution over states at $t-1$ according to the joint distribution over paths ending in states x_{t-1}). Once all the most likely originators have been identified, the most likely terminal state is determined. The overall most likely state path is then found by starting at the most likely terminal state and then follow the path backwards, at each step going from a state to its most likely originator.

Looking at the probabilities in the two-period case and keeping in mind the most likely originator approach, some interesting patterns emerge. Consider a history in which there has been no observed symptoms, i.e. where the signal sequence is $(y_1, y_1, ...) = (r_1, r_1, ...)$. For such a history, the probability of any given path strictly decreases in the length of the history, for all possible paths, *except* that in which neither individual is infected, i.e. the path $(s_1, s_1, ...)$. But this means that for any prior distribution and for any interior values of the parameters (β, σ), the most likely explanation for the absence of symptoms in a sufficiently long sequence of observations, is that neither is in fact infected. To see this, note that the probability assigned to the event that there is no infection in the relationship, conditional on neither individual having shown symptoms for t periods, is given by $P(s_1, s_1, s_1, ...|y_{1:t}) = P(x_0 = s_1)$. This is so because $P(X_{t+1} = s_1|X_t = s_1) = P(y_t = r_1|x_t = s_1) = 1$, which follows since state s_1 is a closed class (i.e. it does not communicate with any other state) and since a healthy individual never shows symptoms. In other words, if state s_1 is the most likely terminal state, then necessarily the most likely state path is $(s_1, s_1, ...)$.[32]

In contrast, the probability of any other path approaches zero, as the number of observations with no symptoms increases. For example, the probability of the path in which both individuals

[30] See e.g. Rabiner (1989) or Russell and Norvig (2003) for more details.

[31] It is worth noting the similarities between the forward algorithm and the Viterbi algorithm. First, where the former algorithm uses the so-called *forward message* $P(X_{t-1}|y_{1:t-1})$, the latter instead uses the message $\max_{\{x_k\}_{k=1}^{t-2}} P(x_1, ..., x_{t-2}, X_{t-1}|y_{1:t-1})$, which are the probabilities of the most likely path to each state x_{t-1}. Second, where the former algorithm sums across states x_{t-1}, the latter instead performs a maximization over states x_{t-1}.

[32] This property is shared with states s_2 and s_3, since individuals cannot recover once they have become infected.

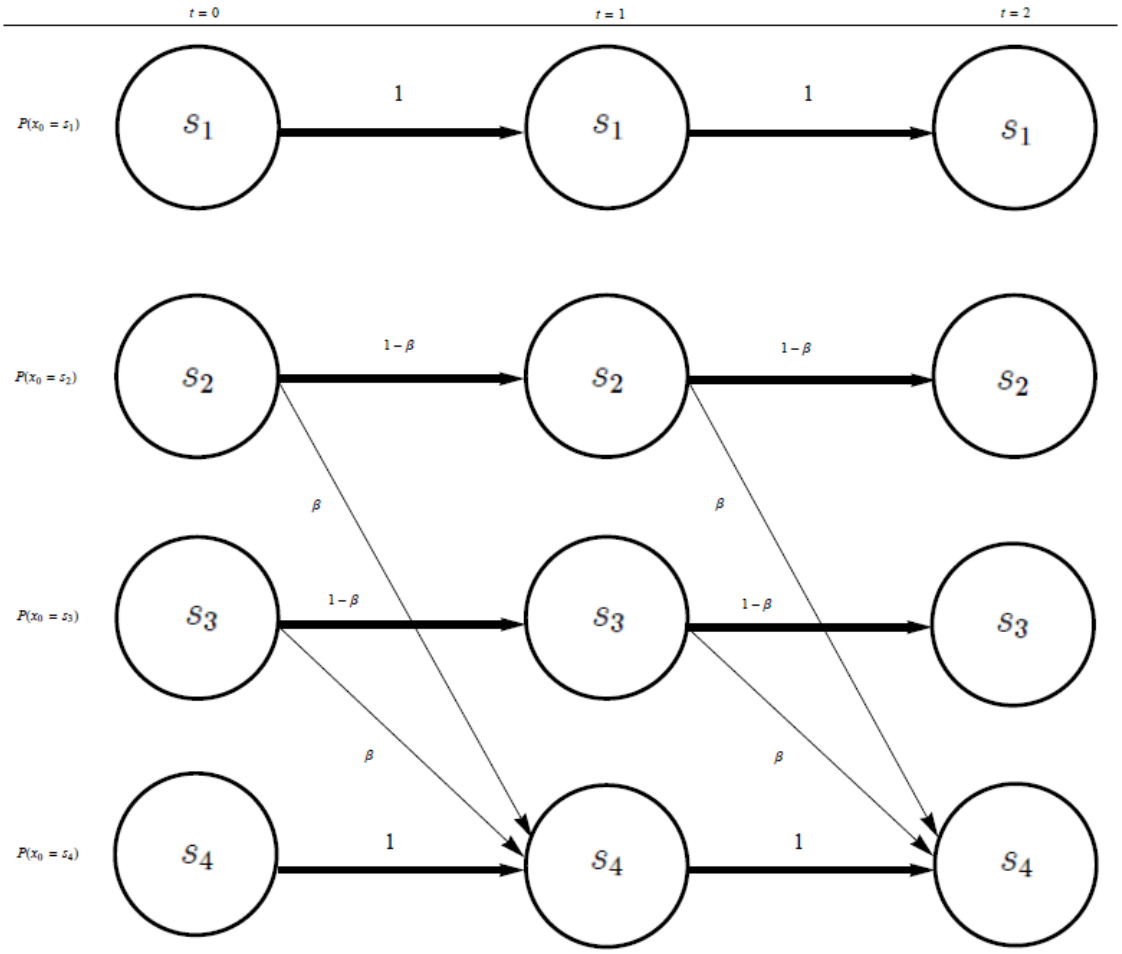

Figure 15: State Transitions and Most Likely Path with Two Observations.

are infected from the outset, is given by $P(s_4, s_4, s_4, ...|y_{1:t}) = P(X_0 = s_4)(1-\sigma)^{2t}$, which vanishes as the number of periods without observed symptoms increases. As the number of possible paths grows, the probability assigned to all these approaches $1 - P(x_0 = s_1)$, even if the probability of any one such path is negligible. In conclusion, when no symptoms are ever observed, the best explanation is that no-one is infected and that the system has always been in state s_1.

It should be emphasized that the Viterbi algorithm returns the state path that has the highest *joint probability* amongst all possible paths. This is distinct from the path returned by the filtering procedure, which yields the posterior distribution across states in a point-wise manner. In other words, a sequence of states that is constructed by including for each period the most likely state in that period (according to the posterior distribution), is not necessarily the best explanation of the observed sequence of signals. To see this, take any prior distribution with $P(x_0 = s_1) > 1/2$ and any observation sequence with $y_t \neq r_1$ for some $= 1, 2, ...$ (i.e. at least one of the individuals displays symptoms at some point in time). Any point-wise most likely path, will necessarily start with $x_0 = s_1$ (because of the chosen prior). But eventually, such a path would have an element $x_u \neq s_1$, since symptoms have been observed. Yet such a path occurs with probability zero, since state s_1 is absorbing, i.e. there cannot be any transition from state s_1 to another state. Another alternative would be to construct a path of the state by selecting at each point in time t the state that is most likely at time t from today's perspective (i.e. with the benefit of hindsight). This process involves so-called *smoothing* and is also subject to the criticism that it is a point-wise procedure that does not determine the joint probability of the entire sequence.[33]

Next, it is interesting to recall the illustrative example considered in the Introduction and illustrated in Figure 1. Suppose that the woman shows symptoms after a few interactions. Under what conditions can one conclude that the woman is the most likely source of infection? Suppose that $(y_1, y_2) = (r_1, r_2)$, i.e. that in the first period, neither is symptomatic but that in the second period, the woman shows symptoms (but the man does not). The possible state paths in which the man is the source of the woman's infection are $\{(s_3, s_3, s_4), (s_3, s_4, s_4)\}$. On the first path, the man infects the woman in period two while on the second path, he infects her in period one. The possible state paths in which the woman is herself infected from the outset, are $\{(s_2, s_2, s_2), (s_2, s_2, s_4), (s_2, s_4, s_4), (s_4, s_4, s_4)\}$. On the first path, the woman is infected at the outset but never infects the man. On paths two and three, the woman infects the man in periods two and one respectively. On the last path, both individuals are infected from the outset.

The probability that the man is the index case, is given by

$$P(x_0 = s_3)\beta\sigma(1-\sigma)^2(2-\beta-\sigma) \tag{78}$$

The probability that the woman is the index case, is given by

$$P(x_0 = s_2)\sigma(1-\sigma)[1-3\beta\sigma+\beta^2\sigma+\beta\sigma^2] \tag{79}$$

The probability that both are infected from the outset, is given by

$$P(x_0 = s_4)\sigma(1-\sigma)^3 \tag{80}$$

It is immediately clear that determining the identity of the index case, is a delicate function of the prior distribution and the parameters of the problem. For the special case in which the prior distribution does not favor either individual, i.e. when $P(x_0 = s_3) = P(x_0 = s_2)$, the fact

[33] Sequences selected on the basis of smoothing may not be feasible paths, given the underlying transition model.

that the woman shows symptoms suggests that she is indeed the index case. But if the man is initially more likely to be infected than the woman (say, because of past exposure), then the man may well be the most likely source of infection. As an example, consider again the case where $\beta = 0.06$ and $\sigma = 0.025$. For $P(x_0 = s_3) = P(x_0 = s_2) = 0.1$ and $P(x_0 = s_4) = 0.01$, the woman is the most likely index case, followed by the man. For $P(x_0 = s_3) = 0.1$, $P(x_0 = s_2) = 0.01$ and $P(x_0 = s_4) = 0.01$, the man is the most likely index case, followed by the woman. In either case, the least likely scenario is that they were both infected from the outset (and thus that there was in fact no disease transmission between the two individuals).

For given initial conditions, it is straightforward to determine the areas in (β, σ)-space where the man, the woman or both are the index case, but these areas vary with the initial conditions in a way that is so unsystematic as to be uninformative.

B.1. Most Likely Paths with Two Observations. The path probabilities when $y_{1:2} = (r_1, r_1)$:

$$
\begin{aligned}
P(s_1, s_1, s_1|y_{1:2}) &= P(x_0 = s_1) && (81)\\
P(s_4, s_4, s_4|y_{1:2}) &= P(x_0 = s_4)(1-\sigma)^4 && (82)\\
P(s_2, s_2, s_2|y_{1:2}) &= P(x_0 = s_2)(1-\beta a_{t-1})^2(1-\sigma)^2 && (83)\\
P(s_3, s_3, s_3|y_{1:2}) &= P(x_0 = s_3)(1-\beta a_{t-1})^2(1-\sigma)^2 && (84)\\
P(s_2, s_2, s_4|y_{1:2}) &= P(x_0 = s_2)(1-\beta a_{t-1})\beta a_{t-1}(1-\sigma)^3 && (85)\\
P(s_3, s_3, s_4|y_{1:2}) &= P(x_0 = s_3)(1-\beta a_{t-1})\beta a_{t-1}(1-\sigma)^3 && (86)\\
P(s_2, s_4, s_4|y_{1:2}) &= P(x_0 = s_2)\beta a_{t-1}(1-\sigma)^4 && (87)\\
P(s_3, s_4, s_4|y_{1:2}) &= P(x_0 = s_3)\beta a_{t-1}(1-\sigma)^4 && (88)
\end{aligned}
$$

Given the observed signals, all other paths have occurred with probability zero.
The path probabilities when $y_{1:2} = (r_1, r_2)$:

$$
\begin{aligned}
P(s_2, s_2, s_2|y_{1:2}) &= P(x_0 = s_2)(1-\beta a_{t-1})^2(1-\sigma)\sigma && (89)\\
P(s_2, s_2, s_4|y_{1:2}) &= P(x_0 = s_2)(1-\beta a_{t-1})\beta a_{t-1}(1-\sigma)^2\sigma && (90)\\
P(s_4, s_4, s_4|y_{1:2}) &= P(x_0 = s_4)(1-\sigma)^3\sigma && (91)\\
P(s_3, s_3, s_4|y_{1:2}) &= P(x_0 = s_3)(1-\beta a_{t-1})\beta a_{t-1}(1-\sigma)^2\sigma && (92)\\
P(s_3, s_4, s_4|y_{1:2}) &= P(x_0 = s_3)\beta a_{t-1}(1-\sigma)^4\sigma && (93)
\end{aligned}
$$

Given the observed signals, all other paths have occurred with probability zero.
The path probabilities when $y_{1:2} = (r_2, r_2)$:

$$
\begin{aligned}
P(s_2, s_2, s_2|y_{1:2}) &= P(x_0 = s_2)(1-\beta a_{t-1})^2\sigma^2 && (94)\\
P(s_2, s_2, s_4|y_{1:2}) &= P(x_0 = s_2)(1-\beta a_{t-1})\beta a_{t-1}(1-\sigma)\sigma^2 && (95)\\
P(s_2, s_4, s_4|y_{1:2}) &= P(x_0 = s_2)\beta a_{t-1}(1-\sigma)^2\sigma^2 && (96)\\
P(s_4, s_4, s_4|y_{1:2}) &= P(x_0 = s_4)(1-\sigma)^2\sigma^2 && (97)
\end{aligned}
$$

Given the observed signals, all other paths have occurred with probability zero.
The path probabilities when $y_{1:2} = (r_2, r_4)$:

$$
\begin{aligned}
P(s_2, s_2, s_4|y_{1:2}) &= P(x_0 = s_2)(1-\beta a_{t-1})\beta a_{t-1}\sigma^3 && (98)\\
P(s_2, s_4, s_4|y_{1:2}) &= P(x_0 = s_2)\beta a_{t-1}(1-\sigma)\sigma^3 && (99)\\
P(s_4, s_4, s_4|y_{1:2}) &= P(x_0 = s_4)(1-\sigma)\sigma^3 && (100)
\end{aligned}
$$

Given the observed signals, all other paths have occurred with probability zero.
The path probabilities when $y_{1:2} = (r_4, r_4)$:

$$P(s_2, s_4, s_4|y_{1:2}) = P(x_0 = s_2)\beta a_{t-1}\sigma^4 \quad (101)$$
$$P(s_3, s_4, s_4|y_{1:2}) = P(x_0 = s_3)\beta a_{t-1}\sigma^4 \quad (102)$$
$$P(s_4, s_4, s_4|y_{1:2}) = P(x_0 = s_4)\sigma^4 \quad (103)$$

C. Exogenous Infection, Imperfect Protection and False Positives

In the treatment in the main text, reality has been simplified by two assumptions, namely that protection is perfect and that initial or sexual disease transmission were the only sources of infection. These simplifications mean that when no symptoms are observed, infection probabilities will converge to the origin. In practice, while sexual transmission is certainly of first order importance in propagating HIV infection, it is by no means the only way that infection can spread. Other potential sources of infection include needle sharing, blood transfusions and accidental needle pricks.[34] Furthermore, even monogamous but serodiscordant couples can cause disease transmission, when infection is imperfect. For these reasons, it is interesting to allow for such additional risk factors and to enrich the model with imperfect protection and with exogenous infection sources. To this end, suppose that in each period, each individual can become independently infected with some probability $\gamma \in [0, 1]$. That is, even under (perfectly safe) protected sex, a susceptible individual will in each period be at risk of infection.[35] To model imperfect protection, recall that under protected sex, the probability of infection for an individual $i = M, W$ is given by $r_{ij}\alpha_S$, where $\alpha_S \in [0, \alpha_R]$ is the probability that protection fails and $\alpha_R \in [0, 1]$ is the probability of transmission when no protection is used. Recall that $\beta \equiv \alpha_R - \alpha_S$ and let $\alpha \equiv \alpha_S > 0$. The risk for individual i from interaction with the partner j, is then given by

$$r_{ij}[a\beta + \alpha] \quad (104)$$

where $a \in \{0, 1\}$ is the joint decision of the couple. Note that with imperfect protection, the risk of infection cannot be entirely eliminated, even disregarding the exogenous source of infection (i.e. $a = \gamma = 0$). The transition matrix $T(a_{t-1})$ under these modifications is given by

$$\begin{pmatrix} 1-\gamma(2-\gamma) & \gamma(1-\gamma) & \gamma(1-\gamma) & \gamma^2 \\ 0 & \omega_1 & 0 & \omega_2 \\ 0 & 0 & \omega_3 & \omega_4 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (105)$$

where

$$\omega_1 \equiv 1 - (\beta a_{t-1} + \alpha) - \gamma(1 - (\beta a_{t-1} + \alpha)) \quad (106)$$
$$\omega_2 \equiv (\beta a_{t-1} + \alpha) + \gamma(1 - (\beta a_{t-1} + \alpha)) \quad (107)$$
$$\omega_3 \equiv 1 - (\beta a_{t-1} + \alpha) - \gamma(1 - (\beta a_{t-1} + \alpha)) \quad (108)$$
$$\omega_4 \equiv (\beta a_{t-1} + \alpha) + \gamma(1 - (\beta a_{t-1} + \alpha)) \quad (109)$$

The extensions to imperfect protection and exogenous infection sources changes the long-term outcomes of the dynamics. First, consider the effect of imperfect protection. Under this assumption, it is clear that unless both individuals are in fact healthy from the outset, they will

[34] Mother-to-child transmission is also a possibility, but one that does not fit neatly into the present analytical framework.

[35] Note that the infection probability is now the probability that the individual be infected by *either* source. Furthermore, once infected, an individual's health status does not differentiate between the two sources of infection, i.e. the individual can now transmit the disease with probability β as in the previous setting.

necessarily end up being both infected. Thus in this case, the only possible long-run outcomes are that beliefs are $(1,0,0,0)$ or $(0,0,0,1)$. Turning to exogenous infection sources, the effects depend on whether these are permanent or transitory. If the couple are permanently exposed to exogenous sources, then they will become infected eventually and the only possible long-run outcome is that beliefs are at $(0,0,0,1)$. If the infection risk is transitory, then it serves as a shock to beliefs (which can change actions and the future evolution of the system), but the analysis is qualitatively unchanged.

In the analysis in the main text, I have assumed that only an infected individual can show symptoms of infection with the disease. That means that upon observing symptoms from a given individual, the beliefs over that individual's health state become degenerate. In practice, symptoms are often ambiguous indicators of the health state of the individual, thereby raising the possibility of false positives. For that reason, I will briefly consider the case in which a symptom can be observed even for a healthy individual. Assume that with probability $\phi \in [0,1]$, a healthy individuals will show symptoms. In this case, the observation process is characterized by the modified emission matrix Q of the form

$$\begin{pmatrix} 1 & (1-\phi)\phi & (1-\phi)\phi & -\phi^2 \\ 1-\sigma & (1-\phi)\sigma & \sigma(1-\phi) & 0 \\ 1-\sigma & (1-\sigma)\phi & \sigma & 0 \\ (1-\sigma)^2 & (1-\sigma)\sigma & \sigma(1-\sigma) & \sigma^2 \end{pmatrix} \tag{110}$$

Under this assumption, beliefs never become degenerate and can, together with the accompanying equilibrium paths of decisions, exhibit rich dynamic behavior. For example, consider an asymptomatic couple who engage in unprotected sex and suppose that one of the individuals suddenly shows symptoms. Such a couple may find it optimal to switch from unprotected to protected sex, in order to protect the asymptomatic individual. If no further symptoms are observed after a period of protected sex, the couple may become sufficiently confident that the symptom was erroneous and therefore switch back to unprotected sex. In that manner, individuals may switch back and forth between unprotected and protected sex indefinitely, because beliefs never become degenerate.

D. Indifference Curves and Iso-Belief Curves in (p_M, p_W)-Space

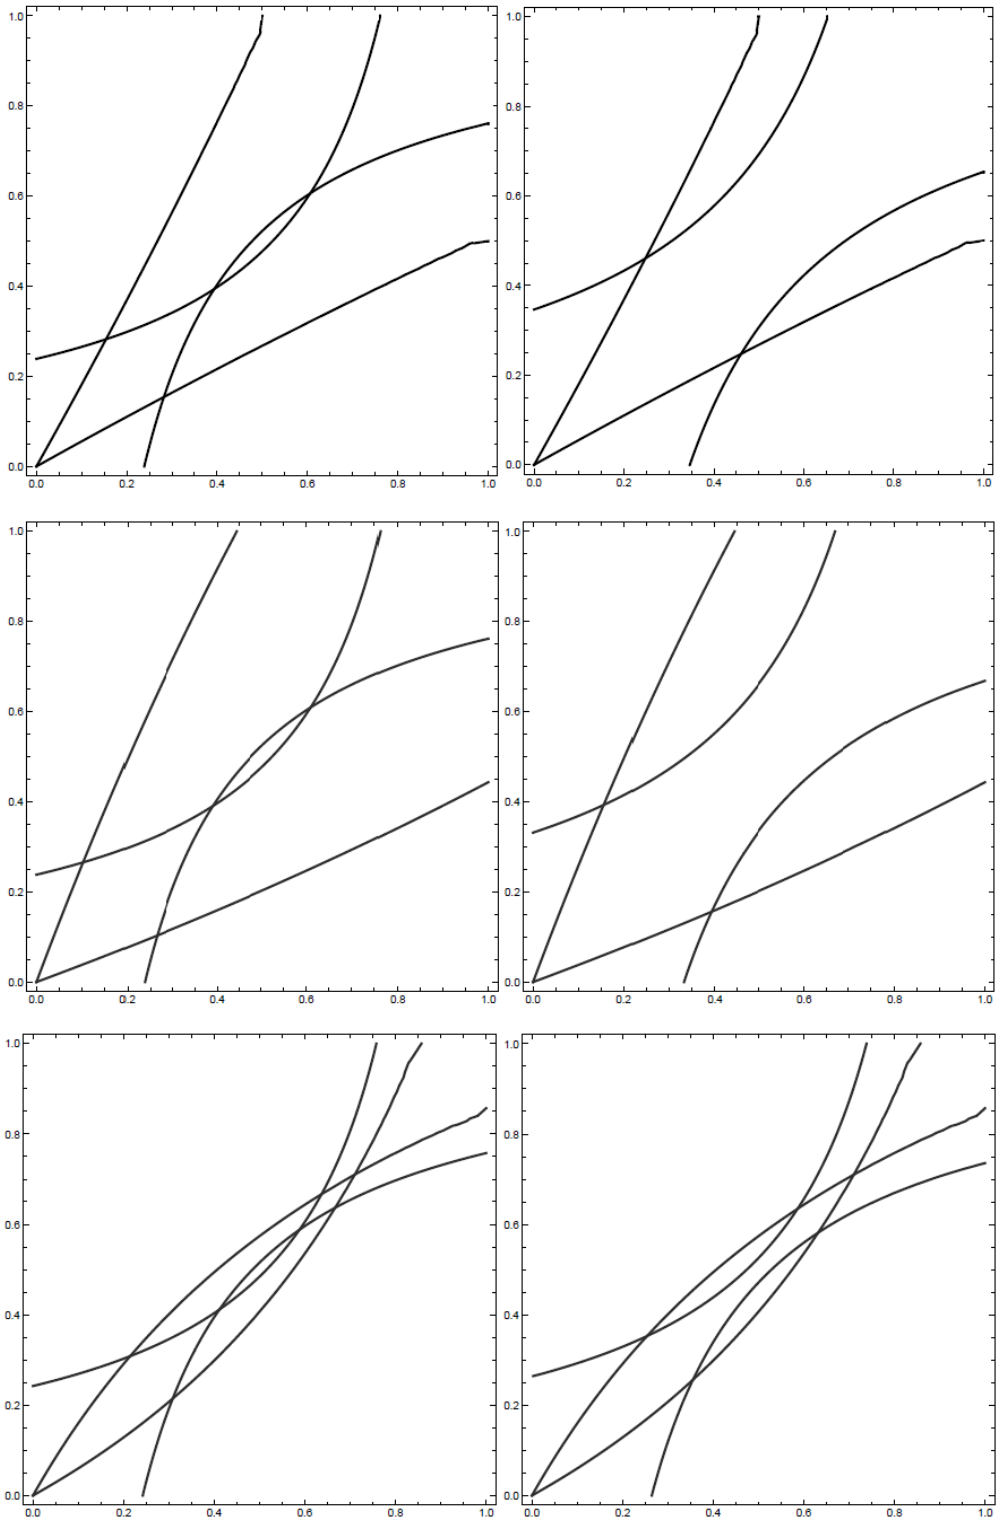

Figure 16: Indifference Curves and Iso-Belief Curves in (p_M, p_W)-Space. Upper panel depicts Case A, middle panel depicts Case B and lower panel depicts Case C. Left-hand side column depicts contiguous cases and right-hand side column depicts non-contiguous cases.

REFERENCES

[1] AKWARA, P. A., N. J. MADISE AND A. HINDE (2003): Perception of Risk of HIV/AIDS and Sexual Behaviour in Kenya, *Journal of Biosocial Science*, 35(3), 385-411.

[2] ANGLEWICZ, P. AND H.-P. KOHLER (2009): Overestimating HIV Infection: The Construction and Accuracy of Subjective Probabilities of HIV Infection in Rural Malawi, *Demographic Research*, 20(6), 65-96.

[3] AULD, C. M. (2003): Choices, Beliefs, and Infectious Disease Dynamics, *Journal of Health Economics*, 22(3), 361-377.

[4] AVERY, C., S. J. HEYMANN AND R. ZECKHAUSER (1995): Risks to Selves, Risks to Others, *American Economic Review: Papers and Proceedings*, 85(2), 61-66.

[5] BACCHETTI, P. AND A. R. MOSS (1989): Incubation Period of AIDS in San Francisco, *Nature*, 338, 251-253.

[6] BELL, D. E. AND R. L. KEENEY (2009): Altruistic Utility Functions for Joint Decisions, in S. J. Brams et al. (eds.), *The Mathematics of Preferences, Choice and Order: Essays in Honor of Peter C. Fishburn*, Studies in Choice and Welfare, Springer0Verlag Berlin Heidelberg.

[7] BOOZER, M. A. AND T. J. PHILIPSON (2000): The Impact of Public Testing for Human Immunodeficiency Virus, *Journal of Human Resources*, 35(3), 419-446.

[8] BRITO, D. L., E. SHESHINSKI AND M. D. INTRILIGATOR (1991): Externalities and Compulsory Vaccinations, *Journal of Public Economics*, 45(1), 69-90.

[9] CENTERS FOR DISEASE CONTROL AND PREVENTION (2010): Monitoring Selected National HIV Prevention and Care Objectives by Using HIV Surveillance Data - United States and 6 U.S. Dependent Areas, *HIV Surveillance Supplemental Report*, 17(3), Part A.

[10] CHEN, F. H. (2009): Modeling the Effect of Information Quality on Risk Behavior Change and the Transmission of Infectious Diseases, *Mathematical Biosciences*, 217(2), 125-133.

[11] CHEN, P.-A. AND D. KEMPE (2008): Altruism, Selfishness, and Spite in Traffic Routing, *Proceedings 9th ACM Conference on Electronic Commerce*.

[12] CHEVRET, S., D. COSTAGLIOLA, J. J. LEFRERE AND A. J. VALLERON (1992): A New Approach to Estimating AIDS Incubation Times: Results in Homosexual Infected Men, *Journal of Epidemiology & Community Health*, 46(6), 582-586.

[13] COOPER, R. AND A. JOHN (1988): Coordinating Coordination Failures in Keynesian Models, *Quarterly Journal of Economics*, 103(3), 441-463.

[14] DELVA, W, F. MENG, R. BEAUCLAIR, N. DEPREZ, M. TEMMERMAN, A. WELTE AND N. HENS (2013): Coital Frequency and Condom Use in Monogamous and Concurrent Sexual Relationships in Cape Town, South Africa, *Journal of the International AIDS Society*, 16:18034.

[15] DELAVANDE, A. AND H.-P. KOHLER (2009): Subjective Expectations in the Context of HIV/AIDS in Malawi, *Demographic Research*, 20(31), 817-874.

[16] Delavande, A. and H.-P. Kohler (2011): HIV/AIDS-Related Expectations and Risky Sexual Behavior in Malawi, http://repository.upenn.edu/psc_working_papers/28.

[17] DeMarzo, P. M., D. Vayanos and J. Zwiebel (2003): Persuasion Bias, Social Influence, and Uni-Dimensional Opinions, *Quarterly Journal of Economics*, 118(3), 909-968.

[18] De Paula, A., G. Shapira and P. E. Todd (2013): How Beliefs About HIV Status Affect Risky Behaviors: Evidence from Malawi, *Journal of Applied Econometrics, forthcoming.*

[19] Dodds J. P., D. E. Mercey, J. V. Parry and A. M. Johnson (2004): Increasing Risk Behaviour and High Levels of Undiagnosed HIV Infection in a Community Sample of Homosexual Men, *Sexually Transmitted Infections*, 80(3), 236-340.

[20] Dow, W. H. and T. Philipson (1996): An Empirical Examination of the Implications of Assortative Matching on the Incidence of HIV, *Journal of Health Economics*, 15(6), 735-749.

[21] Dublin, S., P. S. Rosenberg and J. J. Goedert (1992): Patterns and Predictors of High-Risk Sexual Behavior in Female Partners of HIV-Infected Men with Hemophilia, *AIDS*, 6(5), 475-482.

[22] Duffin, R. (2004): Serostatus, Risk and Responsibility, *Social Research Briefs*, 6.

[23] Eaton, L. A. , T. V. West, D. A. Kenny and S. C. Kalichman (2009): HIV Transmission Risk among HIV Seroconcordant and Serodiscordant Couples: Dyadic Processes of Partner Selection, *AIDS and Behavior*, 13(2), 185-195.

[24] Ellard. J. (2007): Gay Men and Condom Use, Access and Experience, *Queensland Association for Healthy Communities.*

[25] Epstein, H. and M. Morris (2011): Concurrent Partnerships and HIV: An Inconvenient Truth, *Journal of the International AIDS Society*, 14(13).

[26] Received: 18 October 2010Accepted: 15 March 2011Published: 15 March 2011

[27] Fisman, D. N., E. W. Hook 3rd, S. J. Goldie (2013): Estimating the Costs and Benefits of Screening Monogamous, Heterosexual Couples for Unrecognised Infection with Herpes Simplex Virus Type 2, *Sexually Transmitted Infections*, 79(1), 45-52.

[28] Fortenberry, J. D., W. Tu, J. Harezlak, B. P. Katz, and D. P. Orr (2002): Condom Use as a Function of Time in New and Established Adolescent Sexual Relationships, *American Journal of Public Health*, 92(2), 211-213.

[29] Fortenberry, J. D., M. McFarlane, A. Bleakley, S. Bull, M. Fishbein, D. M. Grimley, C. K. Malotte and B. P. Stoner (2002): Relationships of Stigma and Shame to Gonorrhea and HIV Screening, *American Journal of Public Health*, 92(3), 378-381.

[30] Fristedt, B., N. Jain and N. Krylov (2007): Filtering and Prediction: A Primer, *American Mathematical Society.*

[31] Geoffard, P.-Y. and T. Philipson (1996): Rational Epidemics and Their Public Control, *International Economic Review*, 37(3), 603-624.

[32] GEOFFARD, P.-Y. AND T. PHILIPSON (1997): Disease Eradication: Private versus Public Vaccination, *American Economic Review*, 87(1), 222-230.

[33] GERSOVITZ, M. (2011): HIV Testing: Principles and Practice, *World Bank Research Observer*, 26(1), 1-41.

[34] GONG, E. (2015): HIV Testing and Risky Sexual Behaviour, *Economic Journal*, 125(582), 32-60.

[35] GREENWOOD, J., P. KIRCHER, C. SANTOS AND M. TERTILT (2013): An Equilibrium Model of the African HIV/AIDS Epidemic, *mimeo*.

[36] GROENDYKE, C., D. WELCH AND D. R. HUNTER (2011): Bayesian Inference for Contact Networks Given Epidemic Data, *Scandinavian Journal of Statistics*, 38(3), 600-616.

[37] GUIHENNEUC-JOUYAUX, C., S. RICHARDSON AND I. M. LONGINI (2000): Modelling Markers of Disease Progression by a Hidden Markov Process: Application to Characterising CD4 Cell Decline, *Biometrics*, 56(3), 733-741.

[38] HEALTH PROTECTION AGENCY (2012): HIV in the United Kingdom: 2012 Report.

[39] HSU, S.-B. AND Y.-H. HSIEH (2007): On the Role of Asymptomatic Infection in Transmission Dynamics of Infectious Diseases, *Bulletin of Mathematical Biology*, 70(1), 134-55.

[40] JACKSON, C. H., L. D. SHARPLES, S. G. THOMPSON, S. W. DUFFY AND E. COUTO (2003): Multistate Markov Models for Disease Progression with Classification Error, *Journal of the Royal Statistical Society.* Series D (The Statistician), 52(2), 193-209.

[41] JOHNSTON, W. I. (1995): HIV-Negative. How the Uninfected Are Affected by AIDS, *Insight Books*.

[42] KELLER, G. (2007): Passive Learning: A Critique by Example, *Economic Theory*, 33(2), 263-269.

[43] KERWIN, J. T. (2012): "Rational Fatalism": Non-Monotonic Choices in Response to Risk, *mimeo*.

[44] KERWIN, J. T. (2012): The Effect of HIV Infection Risk Beliefs on Sexual Behavior: Scared Straight or Scared to Death?, *mimeo*.

[45] KREMER, M. (1996): Integrating Behavioral Choice into Epidemiological Models of AIDS, *Quarterly Journal of Economics*, 111(2), 549-573.

[46] KU, L., F. L. SONENSTEIN AND J. H. PLECK (1994): The Dynamics of Young Men's Condom Use During and Across Relationships, *Family Planning Perspectives*, 26(6), 246-251.

[47] LUKE, N., S. CLARK AND E. ZULU (2011): The Relationship History Calendar: Improving the Scope and Quality of Data on Youth Sexual Behavior, *Demography*, 48(3), 1151-1176.

[48] MAGDER, L. AND R. BROOKMEYER (1993): Analysis of Infectious Disease Data from Partner Studies with Unknown Source of Infection, *Biometrics*, 49(4), 1110-1116.

[49] MATTHIES, K. AND F. TOXVAERD (2015): Rather Doomed than Uncertain: Risk Attitudes and Sexual Behavior under Asymptomatic Infection, *mimeo*.

[50] Mechoulan, S. (2004): HIV Testing: A Trojan Horse?, *Topics in Economic Analysis and Policy*, 4(1), 1-26.

[51] Moslonka-Lefebvre, M., S. Bonhoeffer and S. Alizon (2012): Weighting for Sex Acts to Understand the Spread of STI on Networks, *Journal of Theoretical Biology*, 311, 46-53.

[52] Muldoon, K. A., P. K. Duff, J. Birungi, M. H. Ngolobe, J. E. Min, R. King, M. Nyonyintono, Y. Chen, K. Shannon, S. Khanakwa and D. Moore (2014): Decisions, Decisions: The Importance of Condom Use Decision Making Among HIV Sero-Discordant Couples in a Prospective Cohort Study in Uganda, *Sexually Transmitted Infections*, 90(5), 408-412.

[53] OByrne, P., A. Bryan and C. Woodyatt (2013): Nondisclosure Prosecutions and HIV Prevention: Results from an Ottawa-Based Gay Men's Sex Survey, *Journal of the Associon of Nurses AIDS Care*, 24(1), 81-87.

[54] ODell, B. L., B. R. S. Rosser, M. H. Miner and S. M. Jacoby (2008): HIV Prevention Altruism and Sexual Risk Behavior in HIV-Positive Men Who Have Sex with Men, *AIDS and Behavior*, 12(5), 713-720.

[55] Oster, E. (2010): Estimating HIV Prevalence and Incidence in Africa from Mortality Data, *B.E. Journal of Economic Analysis & Policy*, 10(1), Article 80.

[56] Ostermann, J., D. S. Brown, A. Muhlbacher, B. Njau and N. Thielman (2015): Would You Test for 5000 Shillings? HIV Risk and Willingness to Accept HIV Testing in Tanzania, *Health Economics Review*, 5(24).

[57] Philipson, T. and W. H. Dow (1998): Infectious Disease Transmission and Infection-Dependent Matching, *Mathematical Biosciences*, 148(2), 161-180.

[58] Philipson, T. and R. A. Posner (1993): Private Choices and Public Health: The AIDS Epidemic in an Economic Perspective, *Harvard University Press*.

[59] Philipson, T. and R. A. Posner (1995): A Theoretical and Empirical Investigation of The Effects of Public Health Subsidies for STD Testing, *Quarterly Journal of Economics*, 110(2), 445-474.

[60] Pinkerton, S. D. and P. R. Abramson (1993): Evaluating the Risks: A Bernoulli Process Model of HIV Infection and Risk Reduction, *Evaluation Review*, 17(5), 504-528.

[61] Rabiner, L. R. (1989): A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition, *Proceedings of the IEEE*, 77(2), 257-285.

[62] Russell, S. and P. Norvig (2003): Artificial Intelligence: A Modern Approach, Second Edition, *Prentice Hall*.

[63] Satten, G. A. and I. M. Longini, Jr. (1996): Markov Chains with Measurement Error: Estimating the 'True' Course of a Marker of the Progression of Human Immunodeficiency Virus Disease, *Journal of the Royal Statistical Society, Series C (Applied Statistics)*, 45(3), 275-309.

[64] Schroeder, K. D. and F. G. Rojas (2002): A Game Theoretical Analysis of Sexually Transmitted Disease Epidemics, *Rationality and Society*, 14(3), 353-383.

[65] SEN, S. (2004). The Behavioral Consequences of HIV Testing: An Experimental Investigation, *Journal of Public Policy & Marketing*, 23(1), 28-42.

[66] SMALL, D. A. AND G. LOEWENSTEIN (2003): Helping a Victim or Helping the Victim: Altruism and Identifiability, *Journal of Risk and Uncertainty*, 26(1), 5-16.

[67] SMITH, L. AND P. SOERENSEN (2008): Rational Social Learning with Random Sampling, *mimeo*.

[68] STONER, B. P., W. L. H. WHITTINGTON, S. O. ARAL, J. P. HUGHES, H. H. HANDSFIELD AND K. K. HOLMES (2003): Avoiding Risky Sex Partners: Perception of Partners' Risks v Partners' Self Reported Risks, *Sexually Transmitted Infections*, 79(3), 197-201.

[69] TEUNIS, P., J. C. M. HEIJNE, F. SUKHRIE, J. VAN EIJKEREN, M. KOOPMANS AND M. KRETZSCHMAR (2013): Infectious Disease Transmission as a Forensic Problem: Who Infected Whom? *Journal of the Royal Society, Interface*, 10(81), 20120955.

[70] TOXVAERD, F. (2014): Recurrent Infection and Externalities in Prevention, *mimeo*.

[71] WILLIAMSON, N. E., J. LIKU, K. MCLOUGHLIN, I. K. NYAMONGOAND F. NAKAYIMA (2006): A Qualitative Study of Condom Use among Married Couples in Kampala, Uganda, *Reproductive Health Matters*, 14(28), 89-98.

[72] WILSON, N. (2010): Antiretroviral Therapy and Demand for HIV Testing: Evidence from Zambia, *mimeo*.