Rather Doomed than Uncertain: Risk Attitudes and Sexual Behavior under Asymptomatic Infection*

Konstantin Matthies† and Flavio Toxvaerd‡

March 2016

Abstract. We analyze the relation between an individual's risk attitude and his willingness to engage in unprotected sexual behavior, when faced with an asymptomatic infectious disease. In such a situation, the individual must not only assess the likely health status of his partner, but must also form beliefs about his own health status, as these two jointly determine the probability of infection transmission. We show that in a high risk environment, increasing an individual's risk aversion counter-intuitively increases his propensity to engage in unprotected sex. The reason for this surprising result is that as risk aversion increases, the certain payoff from getting infected through unprotected sex becomes relatively more attractive than the uncertain payoffs from protected behavior. We also extend the model to analyze the effects of asymmetric information in a population-wide setting.

JEL Classification: I12, D81.

Keywords: Economic epidemiology, risk aversion, sexual behavior, rational fatalism, HIV/AIDS.

1. Introduction

The ongoing HIV/AIDS epidemic is one of the primary public health challenges of our times. To understand how the disease propagates in the population and to design effective programs to manage the epidemic, it is of primary importance to identify the drivers of individuals' behavior and their sexual decision making. Research has found that it is core groups of individuals with particularly high levels of unprotected sexual activity that are key conduits for infection and disease propagation (Shahmanesh et al., 2008). But what characterizes such high activity individuals? Are they simply individuals with a strong intrinsic preference for unprotected sex or are they individuals who are particularly tolerant to the risks involved in transmissive behavior?

In a recent study, Ku et al. (2013) report evidence that

> '[...] if a respondent perceived his partner to be at high risk for HIV, he was *less* likely to use a condom. [...] we posit that selecting a high-risk partner and not using condoms may indicate a general "taste" for risk-taking among those men.'

*We thank Julien Gagnon, Edoardo Gallo, Sönje Reiche and Daniel Sgroi for useful comments on this paper.

†Faculty of Economics, University of Cambridge. Email: krm32@cam.ac.uk.

‡Faculty of Economics, University of Cambridge and CEPR. Address for correspondence: Faculty of Economics, University of Cambridge, Austin Robinson Building, Sidgwick Avenue, Cambridge CB3 9DD, United Kingdom. Email: fmot2@cam.ac.uk.

While this statement is suggestive, it in fact confounds two separate issues. One is the tendency for particular individuals to actively seek out high-risk partners. The second is, given that a high-risk partner has been chosen, the individuals' decision to engage in unprotected sex. To fully disentangle these two issues requires (i) an analysis of individual sexual decision making in the face of uncertainty and (ii) an equilibrium model to account for the sexual matching pattern in the population.[1]

In the same vein, Lammers and van Wijnbergen (2008) state that

> '[...] unsafe sex increases the risk of getting infected, so the more risk-averse one is, the more one should be willing to take precautions to reduce infection risk.'

The suggested link between individuals' risk attitudes and their propensity to self-protect has strong intuitive appeal, but turns out to be overly simplistic. In this paper, we formally characterize how risk aversion influences an individual's willingness to engage in unprotected sexual behavior, controlling for the risk characteristics of the partner. We assume that while unprotected sex is intrinsically desirable (e.g. because protection is costly), it also carries the risk of infection. This is the basic tradeoff faced by the individual in deciding between protected and unprotected sex. A priori, one may think that a more risk averse individual would be less inclined to unprotected sex and thus that, ceteris paribus, those individuals who engage in unprotected sex are likely to be relatively more risk tolerant. This is in the spirit of the work quoted above. For an individual who is almost certainly healthy, this intuition turns out to be correct. This is because protected sex yields a known constant payoff (when protection confers perfect immunity) while unprotected sex leads to uncertainty over the individual's final health status. As risk aversion increases, protected sex becomes even more attractive to the individual, relative to unprotected sex.

The problem with extrapolating from this argument, is that it ignores the possibility that the individual is already (asymptomatically) infected, which may well be the case for a disease like HIV/AIDS. Once this possibility is present, the result can be reversed. To see this, consider the stark case in which unprotected sex leads to infection transmission with certainty. In this case, the expected payoffs from unprotected sex are constant (i.e. known), while protected sex leads to an uncertain outcome (i.e. the basic uncertainty over the individual's initial health status remains unresolved). In this case, a more risk averse individual would value unprotected sex relatively higher than a less risk averse individual. This is not to say that the risk averse individual wants to become infected. Rather, relative to a more risk tolerant individual, he values the relative certainty over outcomes resulting from unprotected sex more than the uncertainty resulting from protected sex. In this paper, we formalize these results and prove that under perfect protection, an increase in risk aversion rotates the individual's indifference curves in such a manner, that individuals facing low risk partners become more inclined towards protected sex, while individuals facing high risk partners become more inclined towards unprotected sex. We show via examples that our results hold for standard utility functions and

[1] In other words, it is possible that high risk individuals are matched with each other simply because no low risk individual would agree to be matched to a high risk individual. Thus a match between high risk individuals may be an equilibrium phenomenon that tells us very little about the underlying risk preferences of the individuals involved.

empirically plausible levels of risk aversion. Thus, although our results may at first be highly surprising, they in fact follow from perfectly mainstream assumptions.

Next, we consider the effects of imperfect protection and show that this possibility leads to unprotected sex becoming even more attractive for more risk averse individuals facing high risk partners. The reason for this, is that imperfections in protection increase the uncertainty over outcomes resulting from protected behavior, while not influencing those resulting from unprotected behavior. We show that when protection is sufficiently imperfect, increasing risk aversion unambiguously increases the propensity to engage in unprotected sex, irrespective of the riskiness of the partner.

Last, we take a different direction and reinterpret the model as one of population-wide matching under asymmetric information and heterogeneous risk attitudes. We argue that in an equilibrium in which the average infection probability of those engaged in unprotected sex is sufficiently high, highly risk averse individuals will tend to be more inclined towards unprotected sex than more risk tolerant individuals.

The reader may wonder whether it is indeed the case that individuals in practice hold diverse beliefs about their own health status. It is conceivable that any individual who harbors doubts about whether they are infected would immediately seek to dispel such doubt and get tested. It turns out that many individuals are (i) uncertain about their own health status and (ii) do not necessarily seek to be tested (see e.g. Delavande and Kohler, 2009 and De Paula et al., 2014). There are many reasons why individuals do not necessarily get tested, amongst them the effects that testing may have on relations with long-term partners (see e.g. discussion in Philipson and Posner, 1993) or the change in legal status that a positive test result may engender (e.g. in the context of disclosure requirements when seeking health insurance). In some situations, individuals may conclude that it is better not to know for sure.[2]

While the literature on sexual decision making and infectious diseases is both voluminous and varied, only a small number of papers are closely related to our setup. Philipson and Posner (1993) were the first to formalize the tradeoffs involved in sexual decision making when individuals are uncertain about their health status and that of their partner(s). They show in a simple static model with a single interaction and protection through condom use, that the riskiness of a potential partner is only one important consideration when choosing whether to avoid transmissive behavior. In particular, they show that for an individual who is sufficiently likely to already be infected, the privately optimal course of action may indeed be to engage in unprotected sex. Under certain conditions, uncertainty about one's own health status can lead to so-called rational fatalism, i.e. the phenomenon that an increase in the perceived riskiness of unprotected sex leads to an increase in the propensity to engage in unprotected sex. Rational fatalism is a central finding in Kremer (1996) and Auld (2002) in models of partner change. In these models, the individual chooses the rate of partner change per period of time. In each period, the individual can therefore have multiple sexual interactions, assumed to yield diminishing marginal utility. For our purposes, it is important to note that rational fatalism in deciding on the marginal exposure to infection is not intrinsically related to risk attitudes, but rather to uncertainty about one's own health status. Indeed, while Kremer (1996) and Auld (2002) implicitly assume that individuals are risk averse, Philipson and Posner

[2] In fact, there is evidence that suggests that not even positive test results for HIV, cause individuals to be certain of HIV positive status (see De Paula et al., 2014).

(1993) confine attention to risk neutral individuals. As long as the perceived riskiness of unprotected sex is allowed to affect an individual's belief about their own health status, rational fatalism can occur. For example, a spike in reported HIV prevalence can on one hand increase the perceived risk of future unprotected sex, as well as increase the belief that one is already infected. If the latter is sufficiently high, an individual may switch from protected to unprotected sex, thus causing a fatalistic response.[3] We should also note that this notion of fatalism is intrinsically different from that studied by Shapiro and Wu (2011).

Taking rational fatalism as given, our work goes beyond this notion and considers the effects that risk attitudes themselves have on sexual decision making under uncertainty, and fatalism in particular. To the best of our knowledge, we are the first to rigorously examine these effects directly. In doing so, we find novel and striking results that have the potential to change the way researchers study sexual decision making.

Last, we should note papers by Greenwood et al. (2014) and Toxvaerd (2014), which study the dynamic interactions between sexual decision making and belief updating in different settings.

In Section 2, we present a model of risky sex with asymptomatic infection. In Section 3, we analyze the effects of risk aversion on the decision to engage in unprotected behavior. In Section 4, we consider the effects of imperfect protection. In Section 5, we consider the extension to matching under asymmetric information. Section 6 concludes and the Appendix contains the main proofs.

2. THE MODEL

Consider an individual who is about to engage in sexual activity with a partner. The individual can choose an action $a \in \{R, S\}$, where $a = R$ denotes unprotected sex and $a = S$ denotes protected sex.[4] Individuals can be in one of two health states, namely healthy or infected. We denote the health state of the individual by $h \in \{H, I\}$, where $h = H$ denotes healthy and $h = I$ denotes infected. The infection is assumed to be asymptomatic, so the individual does not directly observe his health state, nor that of his partner. Instead, we assume that the individual holds a subjective prior belief $p \in [0, 1]$ that he is already infected. Turning to preferences, we assume that the individual is endowed with a strictly increasing and concave utility function $u(x_h^a)$, where x_h^a denotes the level of final 'wealth'. The final wealth level will depend on the health state and the action chosen by the individual. We assume that

$$x_H^R > x_H^S > x_I^R > x_I^S \tag{1}$$

These inequalities mean that while both infection and protection are costly, the cost of infection outweighs the cost of protection. In particular, this implies that if there is no risk of infection transmission (either because both individuals are known to be healthy or both are already infected), the individual will prefer unprotected sex. Furthermore, a healthy individual would rather pay the price of protection, than bearing the cost of

[3]More recent contributions on different aspects of belief formation and fatalism include Sterck (2012) and Kerwin (2012, 2014).

[4]For mnemonic reasons, we use R and S to denote *risky* and *safe* sex, respectively. As will become clear in what follows, the terms *transmissive* and *non-transmissive* behavior are perhaps more appropriate, but our use of language should not cause any misunderstanding.

infection with certainty.

The individual assigns probability $q \in [0,1]$ to the partner being infected and probability $\beta \in [0,1]$ to the infection being transmitted during an unprotected contact between an infected and a healthy individual. Were the individual healthy, he would therefore assign probability $r \equiv \beta q$ to becoming infected as a result of unprotected sex with his partner. We will refer to $r \in [0,1]$ as the risk of infection.

Clearly, the final probability of being infected depends both on the prior belief and on the type of sexual behavior that the individual engages in. For ease of notation, define

$$p^S \equiv p, \quad p^R \equiv p + (1-p)\, r \tag{2}$$

These are simply the ex post probabilities of being infected under protected and unprotected sex, respectively.[5] With this notation, the expected utility of an individual who engages in unprotected sex is given by

$$U^R \equiv \left(1 - p^R\right) u\left(x_H^R\right) + p^R u\left(x_I^R\right) \tag{3}$$

Similarly, the expected utility from protected sex is given by

$$U^S \equiv \left(1 - p^S\right) u\left(x_H^S\right) + p^S u\left(x_I^S\right) \tag{4}$$

The individual's optimal decision depends on the relative magnitudes of U^R and U^S. Indifference between unprotected and protected sex obtains for combinations (r^*, p^*) for which $U^R = U^S$, or

$$(1-p^*)\, u\left(x_H^S\right) + p^* u\left(x_I^S\right) = (1 - [p^* + (1-p^*)\, r^*])\, u\left(x_H^R\right) + [p^* + (1-p^*)\, r^*]\, u\left(x_I^R\right) \tag{5}$$

In (r,p)-space, the indifference curve is given by the function

$$p^*(r) \equiv \frac{(1-r)\left[u\left(x_H^R\right) - u\left(x_I^R\right)\right] - \left[u\left(x_H^S\right) - u\left(x_I^R\right)\right]}{(1-r)\left[u\left(x_H^R\right) - u\left(x_I^R\right)\right] - \left[u\left(x_H^S\right) - u\left(x_I^S\right)\right]} \tag{6}$$

Given a transmission probability r, for $p > p^*(r)$ the individual prefers unprotected sex, while for $p < p^*(r)$, the individual prefers protected sex. Thus the higher the individual's subjective assessment of already being infected is, the more willing will he be to engage in unprotected sex. An example of an indifference curve in (r,p)-space, based on the CRRA utility function, is shown in Figure 1. Note that for combinations (r,p) below the indifference curves, the individual prefers protected sex, while for combinations above it, the individual prefers unprotected sex.[6]

To better appreciate the basic working of the model, consider the left-hand side panel and compare the points a, b and c. A movement from point a to point b constitutes an increase in the risk associated with the partner and will, ceteris paribus, cause the individual to switch from preferring unprotected sex to preferring protected sex. In contrast, a switch from point b to point c constitutes an increase in the probability that the individual is already infected and will, ceteris paribus, cause the individual to switch

[5] Note that in this formulation, protection is perfect and reduces the transmission probability to zero. Later, we will relax this assumption.

[6] The figures in this paper are drawn for the CRRA utility function $u(x) = x^{1-\theta}/(1-\theta)$.

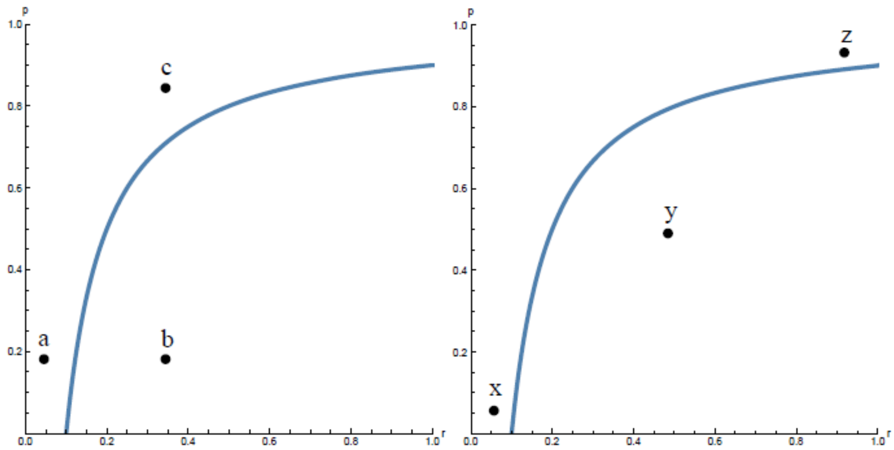

Figure 1: Indifference curves, increasing risks and fatalism.

from preferring protected sex to preferring unprotected sex.

In the right-hand side panel, we illustrate the phenomenon known as rational fatalism. Consider simultaneous increases in both the individual's infection probability and in the partner's risk, e.g. like the movement between the points x, y and z that all lie on the 45°-line. Because the indifference curve is concave in (r, p)-space, the 45°-line can intersect the indifference curve twice. A switch from point x to point y will cause a switch from unprotected to protected sex. But a switch from point y to point z, will cause a switch from protected sex to unprotected sex. This is exactly the behavior known as rational fatalism, i.e. the tendency to respond to a perceived increase in 'risk' by increasing supposedly 'risky behavior'. The different contributions to the literature give different accounts for how a shift like the one from point y to point z can occur. In Philipson and Posner (1993), individuals' infection probabilities are drawn from the same distribution and when the distribution shifts in the sense of first-order stochastic dominance, both the individual's and the partner's infection probabilities increase. In Kremer (1996), individuals commit to a number of exposures. When the overall infection risk increases, either because of increases in aggregate prevalence of because of increases in the transmission probability, the individual 'is moved' from a point like y to a point like z. A similar effect is at work in the model of Kerwin (2012, 2014). Last, a long-term partner's health status may be highly correlated with that of the individual, as shown by Toxvaerd (2014). In this case, once the partner shows symptoms of being infected, the individual may find it more likely to be infected himself, thus prompting a switch to less prevention. This would be consistent with the observation of Ku et al. (2013).

3. Risk Aversion and Sexual Behavior

We will now take a closer look at how the individual's attitudes to risk, determine his decisions vis-à-vis unprotected and protected sex. From inspection of the indifference curve (5), it is trivially true that the decision on protected versus unprotected sex depends on the shape of the utility function, to wit on the individual's risk attitudes.

To trace the effects of varying risk aversion on the individual's decision, we will next determine how the individual's indifference curve moves in (r, p)-space as his coefficient of risk aversion $\theta(x)$ changes.[7]

To trace the shifts in the indifference curves, it is useful to consider the extremes in (r, p)-space. We start by considering what happens as $p \to 0$, i.e. when the individual is virtually certain that he is not already infected. Let $r^*(p)$ denote the value of r that makes the individual indifferent between unprotected and protected sex under prior belief p. We then have the following result:

Proposition 1. *Let $u(x)$ be a well-behaved utility function with local risk aversion coefficient $\theta(x)$. For any $\theta_2(x) > \theta_1(x)\,\forall x$, it is the case that* $\lim_{p\to 0} r_2^*(p) < \lim_{p\to 0} r_1^*(p)$.

This result means that for an individual who can virtually ignore the possibility of already being infected, as his coefficient of risk aversion increases, his willingness to engage in unprotected sex decreases. This results chimes with the common intuition outlined in the introduction. It follows from the observation that for an uninfected individual, protected sex provides a known payoff with certainty (since protection is assumed to be perfect). In contrast, unprotected sex yields an uncertain payoff, because there is uncertainty over whether the individual will become infected. Thus, in this setting, increasing risk aversion unambiguously makes protected sex relatively more attractive. An immediate corollary of this result is as follows:

Corollary 2. *Let $u(x)$ be a well-behaved utility function with local risk aversion coefficient $\theta(x)$. For any $\theta_2(x) > \theta_1(x)\,\forall x$, there exists a range of transmission risks r such that $p_2^*(r) > p_1^*(r)$.*

To understand this result, note that Proposition 1 implies that the indifference curve $p_2^*(r)$ of a more risk averse individual will intersect the r-axis for a smaller r than the indifference curve $p_1^*(r)$. Continuity of equation (6) then implies that there must be a least some range of risks r for which $p_2^*(r)$ lies above $p_1^*(r)$.

Next, we consider another extreme, namely the case where $r \to 1$. In this case, an unprotected encounter will result in the individual ending up infected with certainty (either as a result of being infected by the partner or because the individual was infected at the outset). Our result for this case is a follows:

Proposition 3. *Let $u(x)$ be a well-behaved utility function with local risk aversion coefficient $\theta(x)$. For any $\theta_2(x) > \theta_1(x)\,\forall x$, it is the case that* $\lim_{r\to 1} p_2^*(r) \leq \lim_{r\to 1} p_1^*(r)$.

This result may seem counter-intuitive at first blush. It states that as infection from an unprotected encounter becomes almost certain, an increasingly risk averse individual would become more inclined to engage in unprotected sex! To understand this result, recall our discussion above, comparing uncertain prospects. For an individual that engages in unprotected sex, the outcome is known with certainty when $r \to 1$. That is, the individual knows that in this case, if he engages in unprotected sex, he will end up infected (irrespective of his initial infection status). On the other hand, protected sex (which we have for now assumed leads to perfect protection against any new infection) still leaves the individual with uncertain prospects. In particular, the individual

[7] For concreteness, one can think of $\theta(x)$ as being either the coefficient of absolute risk aversion $A(x) \equiv -u''(x)/u'(x)$ or the coefficient of relative risk aversion $R(x) \equiv -xu''(x)/u'(x)$, but our results do not depend on these particular representations.

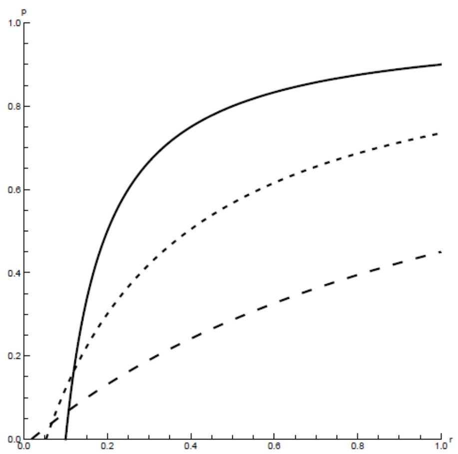

Figure 2: Indifference curves for different degrees of risk aversion.

still assigns the initial probability p to being infected from the outset and probability $(1-p)$ to being healthy. Under protected sex, this uncertainty remains. The upshot of this is that unprotected sex now provides the individual with a known payoff for sure, whereas protected sex means that the individual's prospects remain uncertain.[8] As the individual becomes increasingly risk averse, the uncertainty associated with protected sex becomes increasingly unpalatable, with the individual therefore opting more readily for unprotected sex.

An immediate corollary of this result is as follows:

Corollary 4. *Let $u(x)$ be a well-behaved utility function with local risk aversion coefficient $\theta(x)$. For any $\theta_2(x) > \theta_1(x)\,\forall x$, there exists a range of transmission risks r such that $p_2^*(r) < p_1^*(r)$.*

The rationale for this result is similar to that of the previous Corollary. Figure 2 shows the shifts in the indifference curves for different levels of risk aversion. As can be seen from the figure, an increase in risk aversion causes the indifference curve to pivot clockwise.

It is important to emphasize, that our findings do not rely on the assumption that individuals have unrealistically high levels of risk aversion. To see this clearly, it is instructive to consider an explicit numerical example. Suppose that the individual has constant relative risk aversion and that payoffs are given by $(x_H^R, x_H^S, x_I^R, x_I^S) = (6, 5.5, 0.5, 0)$. For CRRA preferences, Holt and Laury (2002) estimate a coefficient of risk aversion θ in the range $[-1, 1.4]$. As the coefficient of relative risk aversion θ increases from -1 to 1.4, the critical value determining the switch from unprotected to protected sex at risk $r = 1$, decreases from a value around 90% to around 35%. In other words, for individuals in a high risk environment, a reasonable variability in risk aversion can create quite dramatic changes in individuals' willingness to engage in unprotected sexual behavior. This result is quite robust, in that it continues to hold even when the stakes are higher. Consider the alternative example $(x_H^R, x_H^S, x_I^R, x_I^S) = (100, 90, 10, 0)$, where we have amplified the shortfall in utility experienced by an infected individual. In this case, as the coefficient of

[8]This is exactly the sense in which the terms *risky* and *safe* sexual behavior can be misleading.

relative risk aversion θ increases from -1 to 1.4, the critical value determining the switch from unprotected to protected sex, decreases from a value around 98% to around 10%. In fact, the dramatic fall happens even for an individual that faces less extreme risks. In the high-stake example, the shift for an individual facing a more moderate risk, say 50%, would experience a drop from 97% to only 5%. Similar drops will be present for even lower risks, as long as we consider points to the right of the intersection of the indifference curves, which can in turn be located very close to the origin. What these examples show, is that there is nothing extreme or pathological about our setup. The results hold for reasonable preferences and for empirically plausible values of risk aversion.

4. The Effects of Imperfect Protection

In this section, we extend the analysis to allow for the possibility that protected sex does not confer perfect protection against new infection. In particular, we assume that protection is subject to a failure probability $\phi \in [0,1]$. To fix ideas, one can think of ϕ as the probability of condom breakage. While imperfect protection does not alter the expected utility from unprotected sex U^R, the expected utility from protected sex changes to

$$\hat{U}^S \equiv \left(1-p^S\right)\left[(1-\phi)\, u\left(x_H^S\right)+\phi\left[(1-r)\, u\left(x_H^S\right)+r u\left(x_I^S\right)\right]\right]+p^S u\left(x_I^S\right) \qquad (7)$$

To understand this expression, note that with probability $(1-p^S)$, the individual is susceptible to infection. Under protected sex, with probability $(1-\phi)$ the protection works as intended and the individual remains healthy. But with probability ϕ there is a failure in protection, thus exposing the individual to the risk of infection. Whether this risk materializes, depends on the transmission risk r, exactly as is the case for unprotected sex. The indifference curve in this setting is given by the combinations (r,p) such that $U^R = \hat{U}^S$. Solving for the prior p, this gives the new function

$$p^*(r,\phi) \equiv \frac{(1-r)\left[u\left(x_H^R\right)-u\left(x_I^R\right)\right]-\left[\left[(1-\phi r)\, u\left(x_H^S\right)+\phi r u\left(x_I^S\right)\right]-u\left(x_I^R\right)\right]}{(1-r)\left[u\left(x_H^R\right)-u\left(x_I^R\right)\right]-\left[\left[(1-\phi r)\, u\left(x_H^S\right)+\phi r u\left(x_I^S\right)\right]-u\left(x_I^S\right)\right]} \qquad (8)$$

Our first result in this setting is a follows:

Proposition 5. *Let $u(x)$ be a well-behaved utility function with local risk aversion coefficient $\theta(x)$. For any $\theta_2(x) > \theta_1(x)\,\forall x$, it is the case that $\lim_{r\to 1} p_2^*(r,\phi) < \lim_{r\to 1} p_1^*(r,\phi)$.*

This result is equivalent to that under perfect protection, with the levels adjusted for the possibility of protection failure. When infection is the inevitable consequence of unprotected sex, then this behavior yields a known future payoff with certainty. In contrast, protected sex still involves some uncertainty, because the individual has prior probability p of already being asymptomatically infected.

It is more interesting to consider what happens away from the limit, i.e. when $r < 1$. In this case, there is a proper tradeoff in terms of uncertainty, because both unprotected and protected sex offers uncertain prospects. It turns out that this tradeoff is nicely parameterized by the failure probability ϕ. Formally, we have the following result:

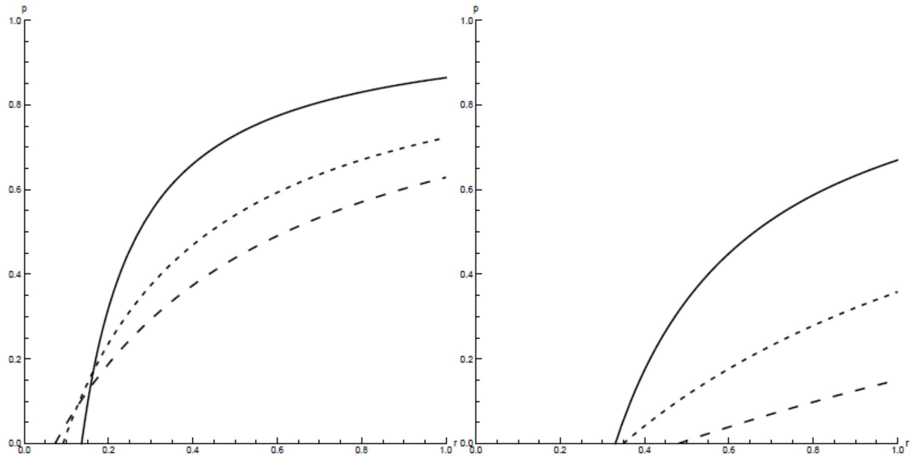

Figure 3: Indifference curves for different degrees of risk aversion and imperfect protection. Left-hand side panel shows low failure probability; right-hand side panel shows high failure probability.

Proposition 6. *Let $u(x)$ be a well-behaved utility function with local risk aversion coefficient $\theta(x)$. For any $\theta_2(x) > \theta_1(x)\,\forall x$, there exists a threshold value $\bar{\phi}$ such that for $\phi \leq \bar{\phi}$, $\lim_{p\to 0} r_2^*(p,\phi) \leq \lim_{p\to 0} r_1^*(p,\phi)$ while for $\phi > \bar{\phi}$, $\lim_{p\to 0} r_2^*(p,\phi) > \lim_{p\to 0} r_1^*(p,\phi)$.*

This result allows us to make a number of interesting observations. First, when the failure probability is sufficiently low (but not necessarily zero), the qualitative features of perfect protection continue hold. That is, as the individual's risk aversion increases, he is more inclined to engage in protected sex if he thinks it likely that he is healthy and more inclined to engage in unprotected sex otherwise.

Second, when the failure probability is sufficiently high, then we can totally order the individual's inclination to engage in unprotected sex by his coefficient of risk aversion. In particular, controlling for the prior p, we have the following:

Proposition 7. *Under imperfect protection with sufficiently high failure probability, the higher the individual's coefficient of risk aversion $\theta(x)$ becomes, the more will he wish to engage in unprotected sex.*

Proof: Let $\mathcal{R}_i$ denote the upper contour set of the indifference curve in (r,p)-space under risk aversion coefficient $\theta_i(x)$, $i = 1, 2$, with $\theta_2(x) > \theta_1(x)\forall x$. This means that if $(r,p) \in \mathcal{R}_i$, then an individual with risk aversion coefficient $\theta_i(x)$ would choose unprotected sex. For sufficiently imperfect protection, the indifference curves do not intersect and thus for $\phi > \bar{\phi}$, $\mathcal{R}_1 \subset \mathcal{R}_2$. This means that $(r,p) \in \mathcal{R}_1 \Rightarrow (r,p) \in \mathcal{R}_2$ and so the higher the level of risk aversion, the higher the propensity to engage in unprotected sex, *irrespective* of prior beliefs p ■

Figure 3 shows the shifts in the indifference curves and the critical values for different levels of risk aversion, for low and high failure probability respectively.

As we have just noted, for $\phi \leq \bar{\phi}$ set inclusion only holds over some ranges of r, because the indifference curves for different risk attitudes intersect. A weaker order than set inclusion is to compare the measures of the risky sets $\mu(\mathcal{R}_1)$ and $\mu(\mathcal{R}_2)$. We have been unable to prove any ranking analytically but have explored the problem numerically under the assumption of perfect protection. From these explorations, it appears that $\mu(\mathcal{R}_1) < \mu(\mathcal{R}_2)$. In other words, as risk aversion increases, there are more pairs (p, r) who prefer to engage in unprotected sex.

5. Private Information and Population Level Matching

In this section, we consider an extension of the model to a population-level setting in which each individual's infection probability is private information. In situations where individuals engage in casual sexual relations with unfamiliar partners, information about an individual's sexual history and past adherence to safe sex practices is bound to be private.[9] To the extent that such information is communicated to new partners at all, the reliability of such information is compromised by the fact that individuals have a considerable incentive to mis-report past sexual behavior. As we will show below, taking such private information into account yields very interesting insights. Schroeder and Rojas (2002) appear to be the only existing authors to consider private information in a disease setting, but their analysis rests on different assumptions than ours.[10]

To discuss population-wide effects of risk aversion, we will re-interpret the model as follows. Rather than considering the choice between protected and unprotected sex in a given encounter, one can alternatively consider an individual's choice between different venues or "markets". The safe option can be thought of as frequenting a venue in which safe sex is the norm, while the risky option can be interpreted as playing the field. To characterize how the population divides itself between the risky and safe options, assume that an individual who chooses the risky option mixes homogeneously within the population of individuals that choose the risky option. For simplicity, we revert to the assumption that protection is perfect.

5.1. The Homogeneous Population Case. To build intuition, we first consider a population which is homogeneous in terms of risk preferences, but heterogeneous in terms of infection probabilities. We will refer to an individual i's probability of being infected as his type. Let the distribution of an individual's possible types be described by the cumulative distribution function F (with associated density function f). Since expected utilities are linear in types, an individual cares only about a potential partner's average type and not about other features of the distribution.

An individual i's strategy is a mapping $a_i : [0,1] \times \Delta[0,1] \rightarrow \{R, S\}$ from the individual's type and the cross-sectional distribution of types to the set of possible decisions, where R stands for unprotected sex and S stands for protected sex.

Let $U_i(p, r) \equiv U^R - U^S$ denote the net utility from unprotected sex. It follows that

$$\frac{\partial U_i(p,r)}{\partial r} = (1-p)\left[u\left(x_I^R\right) - u\left(x_H^R\right)\right] < 0 \quad (9)$$

$$\frac{\partial^2 U_i(p,r)}{\partial r \partial p} = -\left[u\left(x_I^R\right) - u\left(x_H^R\right)\right] > 0 \quad (10)$$

[9]For evidence of this, see e.g. Glynn et al. (2001).

[10]The analysis in Schroeder and Rojas (2002) rests on individuals' ability to signal their infection status to potential partners, but this ability derives from the particular extensive form game that is assumed by the authors. Specifically, it depends on the individuals' ability to commit not to have sex after two rounds of offers for sexual relations.

This means that while an individual clearly has no private incentive to engage in serosorting (i.e. the individual always prefers a partner with as low an infection probability as possible), his incentive to engage in unprotected behavior is increasing in his own type. This means that if an individual of a given type chooses to engage in unprotected sex, then the individual would be willing to do so too for any higher type (controlling for the infection probability of the partner).

Next, we can characterize the identity of the indifferent (or marginal) type engaging in unprotected sex as follows. Suppose that the average of the types of the partner that agree to engage in unprotected sex, is given by some number $q_j \in [0,1]$. Then the marginal type $m_i \in [0,1]$ is implicitly defined by

$$U_i(m_i, q_j) = 0 \tag{11}$$

That is, a type m_i individual facing risk q_j would be exactly indifferent between protected and unprotected sex. The relationship between the average type of the partner and the marginal type of the individual, can be described explicitly as

$$m_i \equiv \max\{0, p^*(r)\} \tag{12}$$

where $p^*(r)$ is defined in equation (6).

Clearly, as q_j increases, so must the marginal type m_i, as one is simply moving up along the individual's indifference curve.

The individual's best response function is of the form

$$a_i(p_i, p_j) = \begin{cases} R & if \quad p_i > m_i \\ S & if \quad p_i \leq m_i \end{cases} \tag{13}$$

where $m_i \in [0,1]$ is calculated using (12).

Next, there is an additional relationship between q_j and m_j, ensuring consistency in the sense that given the marginal type m_j, the average type in the risky pool is indeed q_j. Each value of the marginal type m_j who is indifferent between unprotected and protected sex, induces a truncated distribution of the partner's possible types and hence an associated average infection probability q_j. In particular, from the expectation of a truncated random variable, it follows that

$$q_j = \frac{\int_{m_j}^{1} x f(x) dx}{1 - F(m_j)} \tag{14}$$

Note that this relationship between q_j and m_j is also increasing. To see the intuition for this, note that as m_j increases, the distribution of types who engage in unprotected sex becomes more concentrated on higher types, thereby increasing the average type q_j.

The game played between individuals has two important properties, namely (i) strategic complementarities and (ii) positive spillovers. That the game is one of strategic complementarities, is seen from the fact that the higher others' propensity to engage in unprotected sex, the lower will the average type in the risky pool be. As a consequence, an individual's marginal type will decrease, thereby increasing his propensity to engage in unprotected sex. Similarly, because an individual's gross utility from unprotected sex is decreasing in the average type of the risky pool (and the utility from protected sex is unchanged), there are positive spillovers in the sense that an individual's expected utility is increasing in others' propensity to engage in unprotected sex.

In a symmetric equilibrium, it must be that $m_i = m_j = m^*$ and $q_i = q_j = q^*$ for all i, j. The equilibrium distribution of types that engage in protected and unprotected sex respectively is thus pinned down by a pair (q^*, m^*) that simultaneously satisfy equations (12) and (14). Thus the $F(m^*)$ lowest types will choose to engage in protected sex, while the remaining $(1 - F(m^*))$ highest types will choose to engage in unprotected sex.

Interestingly, since both equations are increasing and generally non-linear, there may be multiple pairs (q^*, m^*) with this property, raising the prospect of multiple equilibria. This result echoes a finding reported in Kremer (1996). In his model, infected and susceptible individuals simultaneously decide whether to join a pool of individuals who engage in unprotected sex. While the infected individuals have a dominant strategy to join, the susceptible individuals play a coordination game between themselves. If enough of them join the risky pool, then the average infection rate in the pool becomes low enough for a susceptible individual to find it optimal to join. Similarly, if only a few susceptible individuals join, then the infection risk in the pool is so high that no susceptible individual would want to join.

The potential equilibrium multiplicity present in this model stems from the strategic complementarities discussed previously. Because there are also positive spillovers, one can Pareto rank these equilibria by the individuals' propensity to engage in unprotected sex.[11] In other words, an equilibrium in which a large proportion of individuals engage in unprotected sex Pareto dominates an equilibrium in which only a small proportion do.

This version of the model can alternatively be interpreted in the context of lemons and the unraveling of markets in the presence of private information. Consider two individuals that, for the sake of argument, are both uninfected. Suppose that the individuals know nothing about each other, except that they both have infection probabilities drawn from the population distribution. Now, since the individuals have no credible means by which they can signal their true types, they must resort to making decisions based on their knowledge of the general population distribution.

To make things interesting, suppose that an individual who is uninfected would not agree to unprotected sex with an individual whose type was the average of the population. The reluctance of this individual to engage in unprotected sex, means that any individual who agrees to unprotected sex must have a strictly positive infection probability. In turn, this means that the relevant average type is not that of the entire population, but rather the average from the truncated distribution obtained by removing the lowest type individuals. The average type of the truncated distribution is necessarily higher than that of the entire distribution. But this means that the two individuals, if they both agree to have unprotected sex, must assign increasingly high probability to each other being infected (otherwise they would not agree to having unprotected sex). The process stops once the first equilibrium (q^*, m^*) is reached. The central observation is that although any two individuals with types sufficiently below m^* would willingly engage in unprotected sex under symmetric information, with private information they would not.

Note that under no circumstances will an individual with type $p < m^*$ agree to engage in unprotected sex. Note also that without private information, such an individual may

[11] The argument follows that in Cooper and John (1988).

well be willing to have unprotected sex with a partner who has the average type of the entire population, say $\widetilde{q}$, but not with the average type when private information is present. In particular, it can happen that for some individual i with $p_i < m^*$,

$$U_i(p_i, \widetilde{q}) > 0 \tag{15}$$

$$U_i(p_i, q^*) < 0 \tag{16}$$

This is exactly the sense in which the presence of private information can reduce social welfare.

The case considered by Kremer (1996) is one in which each individual knows his own infection status but not that of others. In one equilibrium, $m^* = q^* = 1$ and only infected individuals engage in unprotected sex. In the other equilibrium, $m^* = 0$ and $q^* = \widetilde{q}$.

5.2. The Heterogeneous Population Case. We now extend the analysis to a population with heterogeneous risk attitudes. As became clear in previous sections, the individual's decision on whether to engage in unprotected sex depends delicately on his risk attitude. Furthermore, there is in general no way to rank the decisions of individuals in terms of their risk aversion, notwithstanding the assertion of Lammers and Wijnbergen (2008) quoted in the introduction.

Assume that $(p, \theta) \sim G$ on $[0, 1] \times [\underline{\theta}, \overline{\theta}]$ with probability density function g. That is, the population is composed of individuals who are not only heterogeneous in terms of infection probabilities, but also in terms of risk tolerance. Denote the average infection type who engages in unprotected sex by

$$\hat{q} \equiv \frac{\int_{\underline{\theta}}^{\overline{\theta}} \int_{m(\theta)}^{1} x g(x; \theta) dx d\theta}{1 - \int_{\underline{\theta}}^{\overline{\theta}} \int_{0}^{m(\theta)} g(x; \theta) dx d\theta} \tag{17}$$

where the marginal infection type with risk aversion coefficient θ is defined implicitly by the indifference condition

$$U(m(\theta), q) = 0 \tag{18}$$

Recalling equation (6), the marginal infection type is given by[12]

$$m(\theta) \equiv \frac{(1-r)\left[u\left(x_H^R\right) - u\left(x_I^R\right)\right] - \left[u\left(x_H^S\right) - u\left(x_I^R\right)\right]}{(1-r)\left[u\left(x_H^R\right) - u\left(x_I^R\right)\right] - \left[u\left(x_H^S\right) - u\left(x_I^S\right)\right]} \tag{19}$$

We make the following assumptions:

Assumption 1. Assume that $p^*(\beta) < 1$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$.

Assumption 2. Assume that the distribution G is continuous and has full support on $[0, 1] \times [\underline{\theta}, \overline{\theta}]$.

Assumption 1 ensures that even if a match is with an infected individual (which in turn would transmit the disease to an uninfected individual with probability β), there

[12] In fact, the indifferent infection type $m(\theta)$ is also a function of q (through the risk r), but we suppress this dependence in the notation.

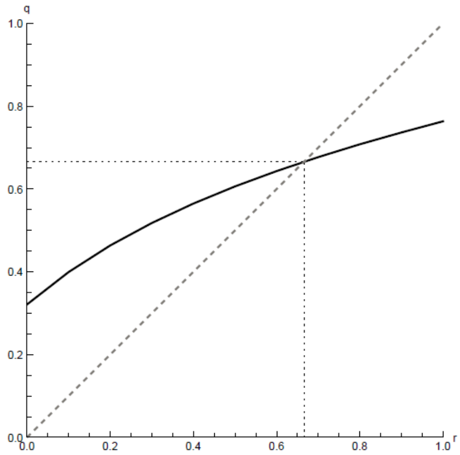

Figure 4: Average infection probability in equilibrium with uniform distribution.

exist types (p, θ) who would prefer unprotected sex. Assumption 2 ensures continuity of the average population type.

Equations (17) and (19) together define a mapping $\hat{q} : [0,1] \to [0,1]$, a fixed point of which is an equilibrium of the model. To appreciate how this works, start by postulating some average type of individual who engages in unprotected sex. This average gives rise to a function of indifference curves (19), one for each risk aversion coefficient $\theta \in [\underline{\theta}, \overline{\theta}]$. Next, given these indifference curves, the average type of individual who engages in unprotected sex is calculated using (17). An equilibrium is reached when the postulated average type who engages in unprotected sex gives rise to the same actual average type.

We can now establish existence of an equilibrium as follows:

Theorem 8. *There exists an equilibrium* $q^* \in [0,1]$ *such that* $\hat{q}(q^*) = q^*$.

Proof: Assume that $q = 0$. Then all individuals would prefer unprotected sex irrespective of type, leading to average type $\hat{q} > 0$. Similarly, assume that $q = 1$. This means that only those individuals who are infected with probability one engage in unprotected sex. But Assumptions 1 and 2 ensure that there are individuals who would choose unprotected sex and who have infection probabilities less than one. Thus $\hat{q} < 1$. The result then follows from the continuity of $\hat{q}(q)$ ■

Figure 4 illustrates the equilibrium in the special case where individuals' types (p, θ) are uniformly distributed on $[0,1] \times [\underline{\theta}, \overline{\theta}]$. The plot shows the 45-degree line and the average type $\hat{q}$. The intersection point of the two curves represents an equilibrium (i.e. a fixed point) of the system.[13]

An important difference between the homogeneous and heterogeneous cases is that in the former, one marginal type completely characterizes the actions of individuals and their type distribution across actions. In particular, any type higher than the marginal type will engage in unprotected sex, while any type lower than the marginal one will abstain from doing so. In the heterogeneous case, there is no one marginal type, but

[13] While the uniform distribution is a tractable assumption to work with, it should be kept in mind that even if the population is initially uniformly distributed, it will not remain so after the first round of contacts. Keeping track of the evolution of the distribution is a significant challenge.

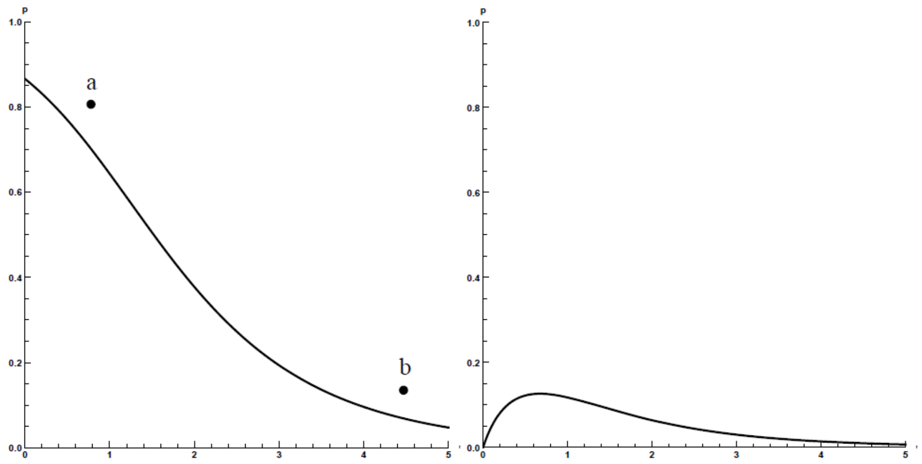

Figure 5: Indifference curves in (θ, p)-space. Left-hand side panel shows high risk environment; right-hand side panel shows low risk environment.

rather a function of marginal types, one for each coefficient of risk aversion. For that reason, it is not necessarily the case it is always the high type individuals who engage in unprotected sex and the low type ones who abstain. Rather, individuals who engage in unprotected sex are those whose combination of infection probability and risk attitude together make unprotected sex the preferred option.

In Figure 5, we display a typical indifference curve of an individual in (θ, p)-space in a high and a low risk environment, respectively. As can be seen in the left-hand side panel, the indifference curve in the high risk environment is downward sloping. This means that for fixed risk, the individual becomes less picky as he becomes more risk averse (i.e. he becomes increasingly willing to engage in unprotected sex). To appreciate the significance of this, consider two points above the curve, such as those denoted by a and b respectively. In this environment (which obtains for a relatively high value of the risk r), both individuals a and b prefer unprotected sex, as they are located above the indifference curve. But note the following important difference. Individual a is very risk tolerant (i.e. he has a low coefficient of risk aversion θ) but is very likely already infected. In contrast, individual b is very risk averse and is unlikely to be infected already; yet he still prefers unprotected sex. Whereas individuals of type a must clearly have more weight in the aggregate than those of type b (for otherwise the average infection type engaging in unprotected sex would be low and it would not be a high risk environment), there may still be a significant number of highly risk averse individuals engaging in unprotected sex who are unlikely to be infected. As can be seen in the right-hand side panel, the indifference curve is not well-behaved in the risk r. As r decreases, the indifference curve becomes hump-shaped. This means that in a low risk environment, it is the most risk averse and the least risk averse individuals who are most inclined towards unprotected sex, while individuals with intermediate levels of risk tolerance that are least inclined towards transmissive behavior.

We now briefly discuss some comparative statics of the model, based on some simple numerical experiments (and using the uniform distribution). We have found the following interesting regularities:

1. Increasing the upper or the lower end of the domain of risk aversion coefficients, decreases the equilibrium average infection type that chooses unprotected sex. In other words, this is an instance in which increasing risk aversion prompts an increase in unprotected behavior.

2. Increasing the cost of protection $(x_R - x_S)$, initially increases the equilibrium average infection type that chooses unprotected sex. In other words, an increase in the cost of protection discourages protection, as would be expected. However, once the protection cost is sufficiently high, further increases prompt a decrease in the average infection type that choose unprotected sex in equilibrium.

3. Increasing the cost of infection $(x_H - x_I)$, increases the equilibrium average infection type that chooses unprotected sex, for relatively high costs of protection. Thus, increasing infection costs discourages unprotected sex, as would be expected. However, for relatively modest protection costs, increasing the cost of infection actually decreases the equilibrium average infection type that chooses unprotected sex. Thus increasing infection costs may encourage more unprotected sex.

It would be interesting to determine the effects of a first-order stochastic dominance shift in the coefficients of risk aversion, for each type in the population. That is, suppose the population generally becomes more risk averse. What would such a shift do to the equilibria of the model and what would the effect on the aggregate amount of unprotected sex be? Unfortunately, it is not possible to make any general statements about this without making strong distributional assumptions. As we have seen in the previous analysis, increasing the risk tolerance of a given individual may lead him to engage in more or less unprotected behavior, depending on his type and the risks he is facing. For that reason, the aggregate effects will depend delicately on the exact population-wide distribution of types and risk attitudes and on the nature of the benchmark equilibrium. What seems clear though, is that it is not unreasonable to expect that a general increase in risk aversion could well lead to an increase in the equilibrium amount of unprotected sex and thus to an aggregate increase in infection. This raises the intriguing possibility that over time, the most risk averse individuals will be over-represented in the set of high type individuals (i.e. those who are most likely to be infected). This would run counter to the widely-held belief that it is a group of high infection type, risk tolerant individuals who drive the spread of the epidemic.

6. CONCLUSION

In this paper, we have demonstrated that individuals' risk attitudes may have non-trivial effects on their propensity to engage in unprotected sexual interactions. In particular, we found that individuals who face sufficiently high risks of infection will become more inclined towards unprotected sex as they become more risk averse. This raises the natural question of how the population-wide distribution of individuals changes over time, a question which we have not addressed in detail.

Based on the analysis in this paper, we are able to address some of the questions raised by the quotes in the introduction. According to Ku et al. (2013), the perception that an individual's partner has a high risk of being infected with HIV, may make the individual less likely to engage in protective behavior. Furthermore, they take this to indicate that such individuals have a preference for unprotected behavior. As we have shown, an individual who believes himself to be infected with high probability may find it privately optimal to have unprotected sex. As shown in Toxvaerd (2014), the health states of long-term partners will tend to be correlated, thereby giving a possible explanation for the first observation. Turning to risk attitudes, we have shown that in a high-risk environment, the willingness to engage in unprotected sex is not necessarily an indication that the individual is risk tolerant, as claimed by Ku et al. (2013) and Lammers and Wijnbergen (2008). In fact, we find the exact opposite to be the case. Specifically, we show that it is the least risk tolerant individuals who are the most likely to engage in transmissive sexual behavior.

It is worth emphasizing that although our results may seem very counter-intuitive at first, they are in fact based on perfectly standard preferences and are not the result of any extreme or unusual assumptions.

This paper is a first step towards a comprehensive analysis of the effects of risk attitudes on sexual decision making under uncertainty. It makes clear that risk attitudes have non-trivial and strongly counter-intuitive effects on individuals' willingness to engage in unprotected sexual behavior, effects that may carry over to population-wide equilibrium patterns of sexual decision making.

A. PROOF OF PROPOSITION 1

The proofs that follow employ a variation of Theorem 1 in Pratt (1964). For simplicity, we shall make use the shorthand notation $u(x,\theta)$. Assume an individual's preferences can be expressed by a utility function $u(x)$ where x is the payoff. Let $\theta(x)$ be the coefficient of local risk aversion (i.e. the Arrow-Pratt measure). An individual is indifferent between exposure and protection if and only if $U^R = U^S$. This will occur at r^*.

First, recall that

$$U^R(\theta_1) = \left(1-p^R\right) u\left(x_H^R,\theta_1\right) + p^R u\left(x_I^R,\theta_1\right) \tag{20}$$

Since $p^R = p + (1-p)r$, it follows that

$$\lim_{p\to 0} U^R(\theta_1) = (1-r)\, u\left(x_H^R,\theta_1\right) + r u\left(x_I^R,\theta_1\right) \tag{21}$$

Similarly, we find that

$$\lim_{p\to 0} U^S(\theta_1) = u\left(x_H^S,\theta_1\right) \tag{22}$$

For indifference, it is hence required that

$$\lim_{p\to 0} U^R(\theta_1) = (1-r_1^*)\, u\left(x_H^R,\theta_1\right) + r_1^* u\left(x_I^R,\theta_1\right) = u\left(x_H^S,\theta_1\right) = \lim_{p\to 0} U^S(\theta_1) \tag{23}$$

By the definition of risk aversion, x_H^S must be the certainty equivalent of the lottery $\left(x_H^R, x_I^R; 1-r_1^*, r_1^*\right)$. Denote this certainty equivalent by $CE_1 \equiv x_H^S$ and note that

$$u\left(CE_1,\theta_1\right) = u\left(x_H^S,\theta_1\right) = (1-r_1^*)\, u\left(x_H^R,\theta_1\right) + r_1^* u\left(x_I^R,\theta_1\right) \tag{24}$$

Now consider an increase in the coefficient of risk aversion to $\theta_2(x) > \theta_1(x)\, \forall x$. Keeping

$r = r_1^*$ fixed, we now have that

$$\lim_{p \to 0} U^R(\theta_2) = (1 - r_1^*)\, u\left(x_H^R, \theta_2\right) + r_1^* u\left(x_I^R, \theta_2\right) \tag{25}$$

From risk aversion, it follows that the certainty equivalent of the gamble $\left(x_H^R, x_I^R; 1 - r_1^*, r_1^*\right)$ must be lower under coefficient $\theta_2(x)$ than under $\theta_1(x)$. Thus

$$CE_2 < CE_1 = x_H^S \tag{26}$$

and it follows that

$$u\left(x_H^S, \theta_2\right) > u\left(CE_2, \theta_2\right) = (1 - r_1^*)\, u\left(x_H^R, \theta_2\right) + r_1^* u\left(x_I^R, \theta_2\right) \tag{27}$$

Again, indifference requires that

$$u\left(x_H^S, \theta_2\right) = (1 - r_2^*)\, u\left(x_H^R, \theta_2\right) + r_2^* u\left(x_I^R, \theta_2\right) \tag{28}$$

Since $x_H^R > x_H^S > x_I^R$, it follows directly that

$$r_2^* < r_1^* \tag{29}$$

This completes the proof ■

B. PROOF OF PROPOSITION 3

It is straightforward to verify that

$$\lim_{r \to 1} U^R(\theta_1) = u\left(x_I^R, \theta_1\right) \tag{30}$$

For indifference it is hence required that,

$$\lim_{r \to 1} U^R(\theta_1) = u\left(x_I^R, \theta_1\right) = (1 - p_1^*)\, u\left(x_H^S, \theta_1\right) + p_1^* u\left(x_I^S, \theta_1\right) = \lim_{r \to 1} U^S(\theta_1) \tag{31}$$

By risk aversion, x_I^R is the certainty equivalent of the lottery $\left(x_H^S, x_I^S; 1 - p_1^*, p_1^*\right)$. Denote this certainty equivalent by $CE_1 = x_I^R$ and note that

$$u\left(CE_1, \theta_1\right) = u\left(x_I^R, \theta_1\right) = (1 - p_1^*)\, u\left(x_H^S, \theta_1\right) + p_1^* u\left(x_I^S, \theta_1\right) \tag{32}$$

Now consider an increase in the coefficient of risk aversion to $\theta_2(x) > \theta_1(x)\, \forall x$. Now, we have that

$$\lim_{r \to 1} U^S(\theta_2) = (1 - p_1^*)\, u\left(x_H^S, \theta_2\right) + p_1^* u\left(x_I^S, \theta_2\right) \tag{33}$$

By risk aversion, it must be that the certainty equivalent of the gamble $\left(x_H^S, x_I^S; 1 - p_1^*, p_1^*\right)$ must be lower under coefficient $\theta_2(x)$ than under $\theta_1(x)$. Thus

$$CE_2 < CE_1 = x_I^R \tag{34}$$

and it follows that

$$u\left(x_I^R, \theta_2\right) > u_i\left(CE_2, \theta_2\right) = \left(1 - p_1^*\right) u\left(x_H^S, \theta_2\right) + p_1^* u\left(x_I^S, \theta_2\right) \tag{35}$$

Indifference requires that

$$u\left(x_I^R, \theta_2\right) = \left(1 - p_2^*\right) u\left(x_H^S, \theta_2\right) + p_2^* u\left(x_I^S, \theta_2\right) \tag{36}$$

Since $x_H^S > x_I^R > x_I^S$, it follows directly that

$$p_2^* < p_1^* \tag{37}$$

This completes the proof ■

C. PROOF OF PROPOSITION 5

Consider first the expression for the expected utility from an unprotected encounter

$$U^R = \left(1 - p^R\right) u\left(x_H^R, \theta_1\right) + p^R u\left(x_I^R, \theta_1\right) \tag{38}$$

Recall $p^R = p + (1-p)r$ and therefore

$$\lim_{r \to 1} U^R = u\left(x_I^R, \theta_1\right) \tag{39}$$

Similarly

$$\lim_{r \to 1} U^S = (1-p)\left[(1-\phi) u\left(x_H^S, \theta_1\right) + \phi u\left(x_I^S, \theta_1\right)\right] + p u\left(x_I^S, \theta_1\right) \tag{40}$$

For indifference it is hence required that, under imperfect protection

$$U^S = \left(1 - p_1^*\right)\left[(1-\phi) u\left(x_H^S, \theta_1\right) + \phi u\left(x_I^S, \theta_1\right)\right] + p_1^* u\left(x_I^S, \theta_1\right) = u\left(x_I^R, \theta_1\right) = U^R \tag{41}$$

We follow a similar approach as in the proof of Proposition 3. Let $\theta_2(x) > \theta_1(x) \, \forall x$. By continuity of $u(x)$ there exists a payoff x', so that $x_H^S > x' > x_I^R > x_I^S$ and

$$(1-\phi) u\left(x_H^S, \theta_1\right) + \phi u\left(x_I^S, \theta_1\right) = u\left(x', \theta_1\right) \tag{42}$$

Substituting this into (41) we obtain

$$\left(1 - p_1^*\right) u\left(x', \theta_1\right) + p_1^* u\left(x_I^S, \theta_1\right) = u\left(x_I^R, \theta_1\right) \tag{43}$$

From concavity of u, it follows that

$$\left(1 - p_1^*\right) u\left(x', \theta_2\right) + p_1^* u\left(x_I^S, \theta_2\right) < u\left(x_I^R, \theta_2\right) \tag{44}$$

Similarly, we have

$$(1-\phi) u\left(x_H^S, \theta_2\right) + \phi u\left(x_I^S, \theta_2\right) < u\left(x', \theta_2\right) \tag{45}$$

It follows, that there exists a $\phi' < \phi$ for which

$$(1-\phi')\,u\left(x_H^S,\theta_2\right)+\phi' u\left(x_I^S,\theta_2\right)=u\left(x',\theta_2\right) \tag{46}$$

Substituting for $u\left(x',\theta_2\right)$ in (44) we obtain

$$\frac{u\left(x_H^S,\theta_2\right)-u\left(x_I^R,\theta_2\right)}{u\left(x_H^S,\theta_2\right)-u\left(x_I^S,\theta_2\right)}<p_1^*+\phi'-p_1^*\phi' \tag{47}$$

Since $\phi' < \phi$ we know that

$$\frac{u\left(x_H^S,\theta_2\right)-u\left(x_I^R,\theta_2\right)}{u\left(x_H^S,\theta_2\right)-u\left(x_I^S,\theta_2\right)}<p_1^*+\phi-p_1^*\phi \tag{48}$$

Indifference requires that $U^R = U^S$ and hence

$$\frac{u\left(x_H^S,\theta_2\right)-u\left(x_I^R,\theta_2\right)}{u\left(x_H^S,\theta_2\right)-u\left(x_I^S,\theta_2\right)}=p_2^*+\phi-p_2^*\phi \tag{49}$$

It follows directly that

$$p_2^* < p_1^* \tag{50}$$

This completes the proof ■

D. Proof of Proposition 6

We will next show that

$$\lim_{p\to 0}\frac{dr^*(p,\phi)}{d\theta}\le 0\ \forall\phi\le\bar{\phi},\quad \lim_{p\to 0}\frac{dr^*(p,\phi)}{d\theta}>0\ \forall\phi>\bar{\phi} \tag{51}$$

We have that

$$\lim_{p\to 0}U^R(\theta_1)=(1-r)\,u\left(x_H^R,\theta_1\right)+ru\left(x_I^R,\theta_1\right) \tag{52}$$

Similarly,

$$\lim_{p\to 0}U^S(\theta_1)=(1-\phi r)\,u\left(x_H^S,\theta_1\right)+\phi ru\left(x_I^S,\theta_1\right) \tag{53}$$

Under imperfect protection, indifference amounts to

$$U^S(\theta_1)=(1-\phi r_1^*)\,u\left(x_H^S,\theta_1\right)+\phi r_1^* u\left(x_I^S,\theta_1\right)=(1-r_1^*)\,u\left(x_H^R,\theta_1\right)+r_1^* u\left(x_I^R,\theta_1\right)=U^R(\theta_1) \tag{54}$$

For completeness, the threshold failure probability is given by

$$\bar{\phi}=\frac{\left(u\left(x_H^S,\theta_2\right)-u\left(x_I^S,\theta_2\right)\right)-\frac{u\left(x_H^S,\theta_1\right)-u\left(x_I^S,\theta_1\right)}{u\left(x_H^R,\theta_1\right)-u\left(x_H^S,\theta_1\right)}\left(u\left(x_H^R,\theta_2\right)-u\left(x_H^S,\theta_2\right)\right)}{\left(u\left(x_H^R,\theta_2\right)-u\left(x_I^R,\theta_2\right)\right)-\frac{u\left(x_H^R,\theta_1\right)-u\left(x_I^R,\theta_1\right)}{u\left(x_H^R,\theta_1\right)-u\left(x_H^S,\theta_1\right)}\left(u\left(x_H^R,\theta_2\right)-u\left(x_H^S,\theta_2\right)\right)} \tag{55}$$

We follow a similar approach as in the previous proofs. Let $\theta_2(x) > \theta_1(x)\,\forall x$. By continuity of $u(x)$, there exists a payoff x', such that $x_H^S > x'' > x_I^R > x_I^S$ and

$$(1-\phi r_1^*)\,u\left(x_H^S,\theta_1\right)+\phi r_1^* u\left(x_I^S,\theta_1\right)=u\left(x'',\theta_1\right) \tag{56}$$

Equation (54) then implies that

$$(1-r_1^*)\,u\left(x_H^R,\theta_1\right)+r_1^* u\left(x_I^R,\theta_1\right)=u\left(x'',\theta_1\right) \tag{57}$$

We know by risk aversion that

$$(1-r_1^*)\,u\left(x_H^R,\theta_2\right)+r_1^* u\left(x_I^R,\theta_2\right)<u\left(x'',\theta_2\right) \tag{58}$$

and also that

$$(1-\phi r_1^*)\,u\left(x_H^S,\theta_2\right)+\phi r_1^* u\left(x_I^S,\theta_2\right)<u\left(x'',\theta_2\right) \tag{59}$$

But (59) implies that there exists a $\phi'' < \phi$ such that

$$(1-\phi'' r_1^*)\,u\left(x_H^S,\theta_2\right)+\phi'' r_1^* u\left(x_I^S,\theta_2\right)=u\left(x'',\theta_2\right) \tag{60}$$

Substituting for $u(x',\theta_2)$ into (58), we obtain the following inequality:

$$\frac{u\left(x_H^R,\theta_2\right)-u\left(x_H^S,\theta_2\right)}{u\left(x_H^S,\theta_2\right)-u\left(x_I^S,\theta_2\right)}<\left[\frac{u\left(x_H^R,\theta_2\right)-u\left(x_I^R,\theta_2\right)}{u\left(x_H^S,\theta_2\right)-u\left(x_I^S,\theta_2\right)}-\phi''\right]r_1^* \tag{61}$$

Recall, that $\phi'' < \phi$. If we let $\phi'' \to \phi$, the right-hand side of (61) decreases. Depending on the relative magnitudes of ϕ'' and ϕ, one of two possibilities arise. Either we have

$$\frac{u\left(x_H^R,\theta_2\right)-u\left(x_H^S,\theta_2\right)}{u\left(x_H^S,\theta_2\right)-u\left(x_I^S,\theta_2\right)}\leq\left[\frac{u\left(x_H^R,\theta_2\right)-u\left(x_I^R,\theta_2\right)}{u\left(x_H^S,\theta_2\right)-u\left(x_I^S,\theta_2\right)}-\phi\right]r_1^* \tag{62}$$

or we have

$$\frac{u\left(x_H^R,\theta_2\right)-u\left(x_H^S,\theta_2\right)}{u\left(x_H^S,\theta_2\right)-u\left(x_I^S,\theta_2\right)}>\left[\frac{u\left(x_H^R,\theta_2\right)-u\left(x_I^R,\theta_2\right)}{u\left(x_H^S,\theta_2\right)-u\left(x_I^S,\theta_2\right)}-\phi\right]r_1^* \tag{63}$$

That is, if ϕ is only slightly larger than ϕ'' then the original inequality remains valid. However, if ϕ is significantly larger than ϕ'', then the inequality is reversed.

Indifference requires the inequalities (62) and (63) to hold with equality for at $r_2^*(0,\phi)$. Hence changing r_1^* to r_2^* in the inequalities will lead to strict equality. Since the right-hand side is strictly increasing in r, it follows that

$$r_2^* \leq r_1^* \tag{64}$$

for (62) and

$$r_2^* > r_1^* \tag{65}$$

for (63). This completes the proof ■

References

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

[2] Delavande, A. and H.-P. Kohler (2009): Subjective Expectations in the Context of HIV/AIDS in Malawi, *Demographic Research*, 20(31), 817-874.

[3] De Paula, A., G. Shapira and P. E. Todd (2014): How Beliefs about HIV Status Affect Risky Behaviors: Evidence from Malawi, *Journal of Applied Econometrics*, 29(6), 944-964.

[4] Cooper, R. and A. John (1988): Coordinating Coordination Failures in Keynesian Models, *Quarterly Journal of Economics*, 103(3), 441-463.

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

[6] Glynn, J. R., M. Carael, B. Auvert, M. Kahindo, J. Chege, R. Musonda, F. Kaona and A. Buve (2001): Why do Young Women Have a Much Higher Prevalence of HIV than Young Men? A Study in Kisumu, Kenya and Ndola, Zambia, *AIDS*, 15(4), 51-60.

[7] Greenwood, J., P. Kircher, C. Santos and M. Tertilt (2014): An Equilibrium Model of the Malawian HIV/AIDS Epidemic, *mimeo*.

[8] Holt, C. A. and S. K. Laury (2002): Risk Aversion and Incentive Effects, *American Economic Review*, 92(5), 1644-1655.

[9] Kerwin, J. T. (2012): “Rational Fatalism”: Non-Monotonic Choices in Response to Risk, *mimeo*.

[10] Kerwin, J. T. (2014): The Effect of HIV Risk Beliefs on Risky Sexual Behavior: Scared Straight or Scared to Death?, *mimeo*.

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

[12] Ku, L, F. L. Sonenstein and J. H. Pleck (2013): The Dynamics of Young Men’s Condom Use During and Across Relationships, *Family Planning Perspectives*, 26(6), 246-251.

[13] Lammers, J. and S. van Wijnbergen (2008): HIV/AIDS, Risk Aversion and Intertemporal Choice, *mimeo*.

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

[15] Pratt, J. W. (1964): Risk Aversion in the Small and in the Large, *Econometrica*, 32(1/2), 122-136.

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

[17] SHAHMANESH, M., V. PATEL, D. MABEY AND F. COWAN (2008): Effectiveness of Interventions for the Prevention of HIV and Other Sexually Transmitted Infections in Female Sex Workers in Resource Poor Setting: A Systematic Review, *Tropical Medicine and International Health*, 13(5), 659-679.

[18] SHAPIRO, J. AND S. WU (2011): Fatalism and Savings, *Journal of Socio-Economics*, 40(5), 645-651.

[19] STERCK, O. (2012): Should Prevention Campaigns Disclose the Transmission Rate of HIV/AIDS? Theory and Application to Burundi, *mimeo.*

[20] TOXVAERD, F. (2014): On the Dynamics of Beliefs and Risky Sexual Behavior, *mimeo.*