NBER WORKING PAPER SERIES

INTEGRATING BEHAVIORAL CHOICE INTO EPIDEMIOLOGICAL MODELS OF AIDS

Michael Kremer

Working Paper 5428

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 1996

I thank Daron Acemoglu for suggestions and Cesaltina Pacheco-Pires, Sergei Severinov, and Robert Shimer for extremely helpful research assistance. This paper is part of NBER's research program in Health Economics. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper 5428
January 1996

INTEGRATING BEHAVIORAL CHOICE INTO EPIDEMIOLOGICAL MODELS OF AIDS

ABSTRACT

Increased HIV risk creates incentives for people with low sexual activity to reduce their activity, but may make high-activity people fatalistic, leading them to reduce their activity only slightly, or actually increase it. If high-activity people reduce their activity by a smaller proportion than low-activity people, the composition of the pool of available partners will worsen, creating positive feedbacks, and possibly multiple steady state levels of prevalence. The timing of public health efforts may affect long-run HIV prevalence.

Michael Kremer
Department of Economics
E52-251C
Massachusetts Institute of Technology
Cambridge, MA 02139
and NBER

Nearly 18 million people have been infected by HIV [WHO, 1995], the majority through heterosexual transmission in developing countries. Prevalence among 30 to 40-year-olds in some districts in Uganda is forty percent [Barnett and Blaikie, 1992], and prevalence among prostitutes in Nairobi reached eighty percent by 1987 [Over and Piot, 1993].[1]

Surveys of sexual activity and epidemiological models suggest that the behavior of a small group of highly sexually active people is critical to the spread of the epidemic [Hethcote and Yorke, 1984; Over and Piot, 1993]. However, most epidemiological models treat behavior as independent of prevalence. This may be in part because there is little data on how prevalence affects the rate of partner change, and in part because the available evidence suggests that people have reacted differently to the epidemic. Although some people have changed their behavior in response to prevalence [McKusick, et. al., 1985; Ahituv, Hotz, and Philipson, 1993], there is anecdotal evidence of fatalism among some IV drug users and homosexuals in developed countries, and prostitution has continued at high levels in parts of Africa and Asia.[2]

This paper integrates formal analyses of behavioral choice and epidemiological dynamics in heterogeneous populations. Unsurprisingly, increases in the probability of infection from an additional partner will create incentives for people with low sexual activity to further reduce their activity. In homogeneous, low-prevalence populations, this will lead to a unique stable steady-state prevalence of the disease. If this behavioral response is elastic,

---

[1]Cited in Piot et al., [1990].

[2]This does not seem to be due to ignorance about the disease, since surveys indicate people are knowledgeable about the disease [Barnett and Blaikie, 1992].

the introduction of imperfect vaccines will increase prevalence and impose negative externalities.

On the other hand, those with sufficiently high sexual activity may only reduce their activity slightly, or even increase it, in response to increased prevalence in the pool of available partners. This is because increased prevalence will make highly-sexually active people so likely to be infected by their infra-marginal partners that their marginal probability of infection from an additional partner changes little or actually declines. Barnett and Blaikie [1992] report anecdotal evidence of fatalism in high-prevalence regions of Uganda.

If increases in prevalence spur smaller proportional reductions in activity by high-activity people than by low-activity people, the composition of the pool of available partnerships will worsen. For example, suppose that people choose sexual partners randomly from a bar, and that half the population goes to the bar and picks a new partner ten times a year, and the other half goes five times a year. Given that high-activity people will be over-represented in the bar on any given night, the probability of picking a high-activity partner will be 10/15ths. If increases in prevalence cause each group to reduce its rate of partner change by one partner a year, to nine and four respectively, the chance of meeting a high-activity person will increase to 9/13ths. The increased probability that a randomly selected partnership will be with an infected person will lead to further reductions in the rate of partner change, creating a positive feedback. If this positive feedback is strong enough, there may be multiple equilibria.

Those with sufficiently high enough activity may respond to increased prevalence by

actually increasing their activity. In this case, there may be one steady-state equilibrium in which some people have very few partners and others have very many, so there is a large probability that a randomly selected partner will be infected; and another steady-state equilibrium in which all people have a moderate number of sexual partners, so there is a small probability that a randomly selected partner will be infected. The timing of public health efforts may influence long-run prevalence. High-prevalence areas in which prostitution has continued, such as parts of Africa and Thailand, may be trapped in the former equilibrium.

This analysis differs both from traditional epidemiological analyses, which take behavior as independent of prevalence, and from the few attempts to introduce behavioral considerations into epidemiology, which do not formally model decisions about the rate of partner change given the composition of the pool of available partners. Models without behavioral response typically generate a unique stable steady state. Analyses in which behavior is an ad hoc declining function of prevalence typically imply that anti-AIDS policies, such as prevention efforts or imperfect vaccines, will be counteracted by behavioral response, at least partially, and possibly more than fully [Philipson and Posner, 1993; Hadeler and Castillo-Chavez, 1994]. If these analyses implicitly or explicitly assume that behavior is a function of lagged prevalence, they will generate cycles in prevalence [Avery, Heymann, and Zeckhauser, 1994; Brauer, Castillo-Chavez, and Velasco-Hernandez, 1994].[3] In contrast,

---

[3]Brauer, Castillo-Chavez, and Velasco-Hernandez [1994] generate multiple steady states and cycles in models in which the rate at which people join the highly sexually active population depends on prevalence. In these models, the change in the rate of partner change, rather than the rate of partner change itself, depends on prevalence, and therefore the rate of partner change

this analysis, which formally models the way in which decisions about the rate of partner change depend on the composition of the pool of available partners, suggests that there may be positive feedbacks in behavior, and possibly even multiple steady-states. This implies that behavioral response may sometimes increase the effectiveness of prevention policies, and that the timing of prevention efforts may determine whether society winds up in a favorable or unfavorable steady-state.

Simultaneous formal modeling of behavior and epidemiological dynamics has the additional advantages of making clear the circumstances under which increased prevalence can make people fatalistic, and thus increase activity; of clarifying the conditions under which imperfect vaccines can increase steady-state prevalence; and of making clear the welfare implications of the introduction of imperfect vaccines in such circumstances. It also makes clear that cycles are not a generic consequence of behavioral response, but instead are generated by the assumption that behavior responds to prevalence with a lag.[4]

Surveys in developed countries indicate that the main ways people claim to have responded to the AIDS epidemic are by using condoms and reducing their number of partners [Wellings, et al., 1994]. Since condom use has been examined by other authors [Ahituv, Hotz, and Philipson, 1993], and this paper focuses on Africa, where condom use is not extensive, this paper focuses on changes on the number of partners. Similarly, given that this

---

depends on lagged prevalence.

[4] While people might use a rule of thumb such as responding to the number of AIDS cases or deaths early in the epidemic, it does not seem likely that they would continue to do so if a regular cycle were established.

paper is concerned primarily with Africa, I will assume that HIV testing is unavailable and that people do not know whether they are infected. Other types of response, such as HIV testing, adoption of (hypothetical) vaccines, and changes in partner selection, are explored by Philipson and Posner [1993], Geoffard and Philipson [1995], Philipson [1995], Kremer [1994], and Boozer and Philipson [1995]. Although this paper focuses on AIDS, the model should also be applicable to other sexually transmitted diseases.

The paper is organized as follows. Section I integrates simple models of behavioral choice and epidemiological dynamics in homogeneous populations, and examines the impact of imperfect vaccines. Section II extends the model to randomly-mixing, heterogeneous populations. It shows that increases in prevalence may lead to behavioral changes which worsen the composition of the pool of available partners, creating a positive feedback. Section III argues that in high prevalence populations, there may be one equilibrium in which differences in sexual activity are extreme, and another in which they are muted. Section IV agues that if high-activity people are disproportionately likely to match with each other, the scope for positive feedbacks will be increased.

I. Homogeneous Populations

This section modifies a standard epidemiological model to allow activity to depend on prevalence. It shows that low-activity people will reduce activity in response to increased prevalence, and that this will tend to produce a unique stable steady-state prevalence. If

behavioral response is elastic, imperfect vaccines may increase prevalence and create negative externalities. High-activity people, on the other hand, may increase activity in response to increased prevalence.

Under the standard susceptible - infected (SI) epidemiological model, as set forth by Anderson and May [1992], people are born and die with Poisson hazard rate δ, the transmission rate is denoted β, and the rate of partner change is denoted i.[5] (Epidemiologists typically treat i as a parameter, but in this paper, i will depend on prevalence, Y.) The change in prevalence is equal to the number of uninfected people who have sex with infected people and become infected, minus the number of infected people who die:

$$\dot{Y} = i(Y)\beta Y(1-Y) - \delta Y. \tag{1}$$

The expected number of secondary infections that a single infected person will cause in an otherwise uninfected population is denoted R_0. If $R_0 < 1$, the disease will die out. In this model, R_0 equals i(0), the rate of partner change if prevalence is zero, times β, the transmission rate, times $1/\delta$, the expected lifespan. To find the stable steady-state level of prevalence for $R_0 > 1$, set $dY/dt = 0$ and divide by Y to obtain $Y^* = 1-\delta/i(Y^*)\beta$.

I will assume that individuals choose their rate of partner change, i, to maximize the utility function u(i)-p(i), where p(i) denotes the probability of becoming infected and $u'' < 0$,

---

[5]This model does not allow for the fact that infected people have a higher death rate. While higher death rates among infected people will reduce steady-state prevalence, there is little reason to believe they would change the qualitative results of this paper. Regardless of the death rate, increases in the probability of infection from an additional partner will create incentives for people to reduce their activity, and greater proportional reductions in sexual activity by low-activity people will worsen the composition of the pool of available partners. Kremer [1995] reports simulations in which infected people have a higher death rate.

so there is diminishing marginal utility from additional partners.[6] Note that this does not allow for altruism. In order to keep the model tractable, I will assume that people cannot allow their rate of partner change to depend on their age.[7] I will also assume that there is some maximum rate of partner change, i_{max}, which can be thought of as either a physical maximum, or the rate of partner change at which $u'(i) = 0$, so that people would not want more partners even in the absence of HIV risk.[8]

A. Choice of Rate of Partner Change for Low-Activity People

If β/δ and i_{max} are small enough, the probability of infection will be approximately linear in the number of partners. This implies that the rate of partner change will decrease in the marginal probability of infection from an additional partner, βY. In this case, there will be a unique fixed point of $Y^* = 1\text{-}\delta/i(Y^*)\beta$, since Y^* is continuous and non-decreasing in i, i

---

[6]Note that I assume utility depends on the chance of ever being infected and not on the date of infection. In a high-prevalence population, a lower rate of partner change will not only reduce the chance of ever being infected, but will also significantly delay the expected date of infection in the case in which the person is infected. However, since people do not die of AIDS for many years after being infected, people may not place much weight on whether they die in fifteen years versus twenty years, compared to the weight they place on whether or not they die of AIDS at all.

[7]Theoretically, as people get older, and have more previous partners, so they are more likely to already be infected, they would like to increase their rate of partner change. This effect is not likely to be important in low-prevalence populations, but could be important in high-prevalence populations. However, in practice, it is likely to be outweighed by changing preferences for the rate of partner change with age.

[8]I assume that there is no intertemporal substitutability in demand for partners. If there were, then in the early stages of an epidemic, people would have more partners, and hence the disease would spread faster.

is continuous and decreasing in Y^*, and Y^* is bounded in [0,1]. The steady state will be stable, since prevalence increases below the steady state and falls above it.

Reductions in β, the transmission rate, will lead to an increase in i, the desired number of partners. If the elasticity of the rate of partner change to the marginal probability of infection from an additional partner is one, any reduction in β will be exactly counter-balanced by increases in i, leaving $Y^*=1-\delta/(i\beta)$ unchanged. Reductions in the transmission rate will increase steady-state prevalence if the elasticity is greater than one, and will reduce prevalence if the elasticity is less than one. (This result arises because steady-state prevalence depends on the product of the rate of partner change and the transmission rate. It is not analogous to the familiar partial equilibrium result that total expenditure on a good will increase in response to a fall in price if demand is elastic. In this case, the "price" of an increase in activity in terms of the additional probability of infection depends on the activity levels of other agents.)

Fears that partially effective vaccines will lead to increases in activity, and could raise prevalence, [Hadeler and Castillo-Chavez, 1994], led to an indefinite hold on U.S. testing of gp120, an HIV vaccine that might be 30 percent effective [Green, 1995]. Whereas epidemiologists and public health officials typically evaluate policies by their effect on prevalence, economists have argued that since people have preferences over the desired number of partners, as well as the risk of infection, policies should be evaluated using a more comprehensive index of social welfare [Philipson and Posner, 1993]. This analysis suggests that holding constant income and other variables, an individual's *ex ante* welfare is determined

by the tradeoff he or she faces between avoiding infection and the desired number of partners. This, in turn, is determined by βY, the probability of infection from an additional partner.

Surprisingly, the availability of an imperfect vaccine could reduce welfare, as well as increase prevalence. If the elasticity of the rate of partner change to the marginal probability of infection from an additional partner is greater than one, an individual who gains access to a partially effective vaccine will wind up with a higher risk of infection than without the vaccine. Although the individual benefits, others are made worse off, because prevalence in the pool of available partners will increase. Adoption of an imperfectly effective vaccine could not cause the number of partners to increase so much that βY increased, because people would not be willing to have more partners if the probability of infection from an additional partner increased. However, the combined costs of the increased prevalence, plus the expense and side effects of the vaccine, could outweigh the benefits of a reduced risk of infection per partner and so introduction of an imperfect vaccine could make everybody worse off. This analysis suggests that one goal for empirical work should be to determine whether the elasticity of behavioral response to the probability of infection is less than or greater than one.

Note that the holder of a patent on an imperfect vaccine will have incentives to adjust the price of the vaccine in exactly whichever manner is counter to what is socially optimal. If behavioral response is elastic, the patent holder will subsidize the vaccine relative to the monopoly price for prevention of a non-contagious disease, because by doing so the

monopolist will increase the incentive for other people to purchase the vaccine.[9] On the other hand, if demand for the vaccine is inelastic, so that people who adopt the vaccine will have lower prevalence, monopolists will charge more for the vaccine than they would for a similar preventative for a non-contagious disease, because a sale to one person will reduce demand for the vaccine by other people.

B. High-Activity People

The analysis above applies only to those with low activity, because it assumed that the probability of infection could be approximated as linear in the rate of partner change. More generally, the probability that someone will be infected before dying of other causes, denoted p(i), is $i\beta Y/(\delta + i\beta Y)$. Thus $\partial p/\partial i = \delta\beta Y/(\delta + i\beta Y)^2$ and

$$\frac{\partial^2 p}{\partial i \partial Y} = \frac{\delta\beta[\delta - i\beta Y]}{\delta + i\beta Y}. \tag{2}$$

This cross-derivative represents the effect of an increase in prevalence on the marginal risk of infection from an additional partner. Note that as long as this cross-derivative is positive, it is decreasing in i, so that increases in prevalence will have a smaller effect on the marginal probability of infection for people with higher activity.

Note from equation (2) that for those with enough partners that $i > \delta/\beta Y$, increases in prevalence in the pool of available partners will actually reduce the marginal risk of infection

[9]Geoffard and Philipson [1995a] independently make this point and explore it in more detail.

from an additional partner. This rate of partner change corresponds to a 1/2 chance of being infected, since $p(i) = i\beta Y/(\delta + i\beta Y)$.[10] Increases in prevalence in the pool of available partners may thus cause people with more than a fifty percent chance of becoming infected over their lifetime to increase their rate of partner change.[11]

Figure I illustrates that for those with $i>\delta/\beta Y$, the locally optimal number of partners will increase in response to increases in prevalence in the pool of available partners. The first order condition for the desired number of partners is that the marginal cost, in increased probability of infection, from an additional partner equals the marginal utility from an additional partner. An increase in prevalence will increase the marginal probability of infection for the first few partners, but will eventually reduce the marginal probability of infection, and thus will cause the locally optimal number of partners to increase. The globally optimal number of partners may decrease if the marginal utility curve crosses the new marginal probability of infection curve at a lower number of partners, as illustrated with the dotted line. Note that this is most likely if the marginal utility curve contains flat portions, which correspond to an elastic behavioral response to prevalence.[12]

---

[10]As shown in earlier versions of this paper, under a model with a fixed lifetime, rather than a constant probability of death, the cutoff probability of infection is 1-1/e, or approximately 0.63, rather than 1/2. This figure may be more realistic.

[11]If people find out about the increase in prevalence after having already had i^* partners, they will increase their desired number of partners no matter what their elasticity of behavioral response.

[12]Theoretically, increasing activity at high-prevalence could lead to multiple non-trivial steady-state level of prevalence, but this would not be empirically relevant if populations were homogeneous, because there is no national population in which prevalence is greater than the 50 percent required for these counterintuitive effects.

It may seem difficult to believe that people could actually increase their activity in response to an increase in prevalence. However, this model suggests that it would be a mistake to assume that intuition developed from experience with people who have a small chance of becoming infected with HIV can be used to generalize about people with high prevalence. Most non-IV drug-using heterosexuals in rich countries will have a probability of infection far too low to induce fatalism, even if they are relatively high-activity. Thus, this group is likely to reduce their activity substantially in response to increased prevalence. Barnett and Blaikie [1992] find that some people in Uganda's high-prevalence Rakai district have become fatalistic, and it is clear that given their bleak alternatives, people are continuing to enter prostitution despite the high prevalence of HIV. In Rakai, prevalence is 30 percent among 30- to 40-year-olds, so prevalence among high-activity people is presumably well above the 50 percent threshold at which high-activity people may increase their activity in response to increased prevalence. Even in lower prevalence populations, high-activity people may have enough partners that increased prevalence in the pool of available partners would not lead to significantly reduced activity.[13]

Although I have argued that increased prevalence may lead high-activity people to increase their activity, the argument in Section II does not require that high-activity people actually increase activity in response to increased prevalence, but merely that they reduce their activity by a smaller proportion than low-activity people. The argument in Section III

---

[13]If some people are inherently more prone to infection than others, this tendency towards decreasing marginal cost will be exacerbated.

Figure I
Desired Number of Partners as a Function of Prevalence

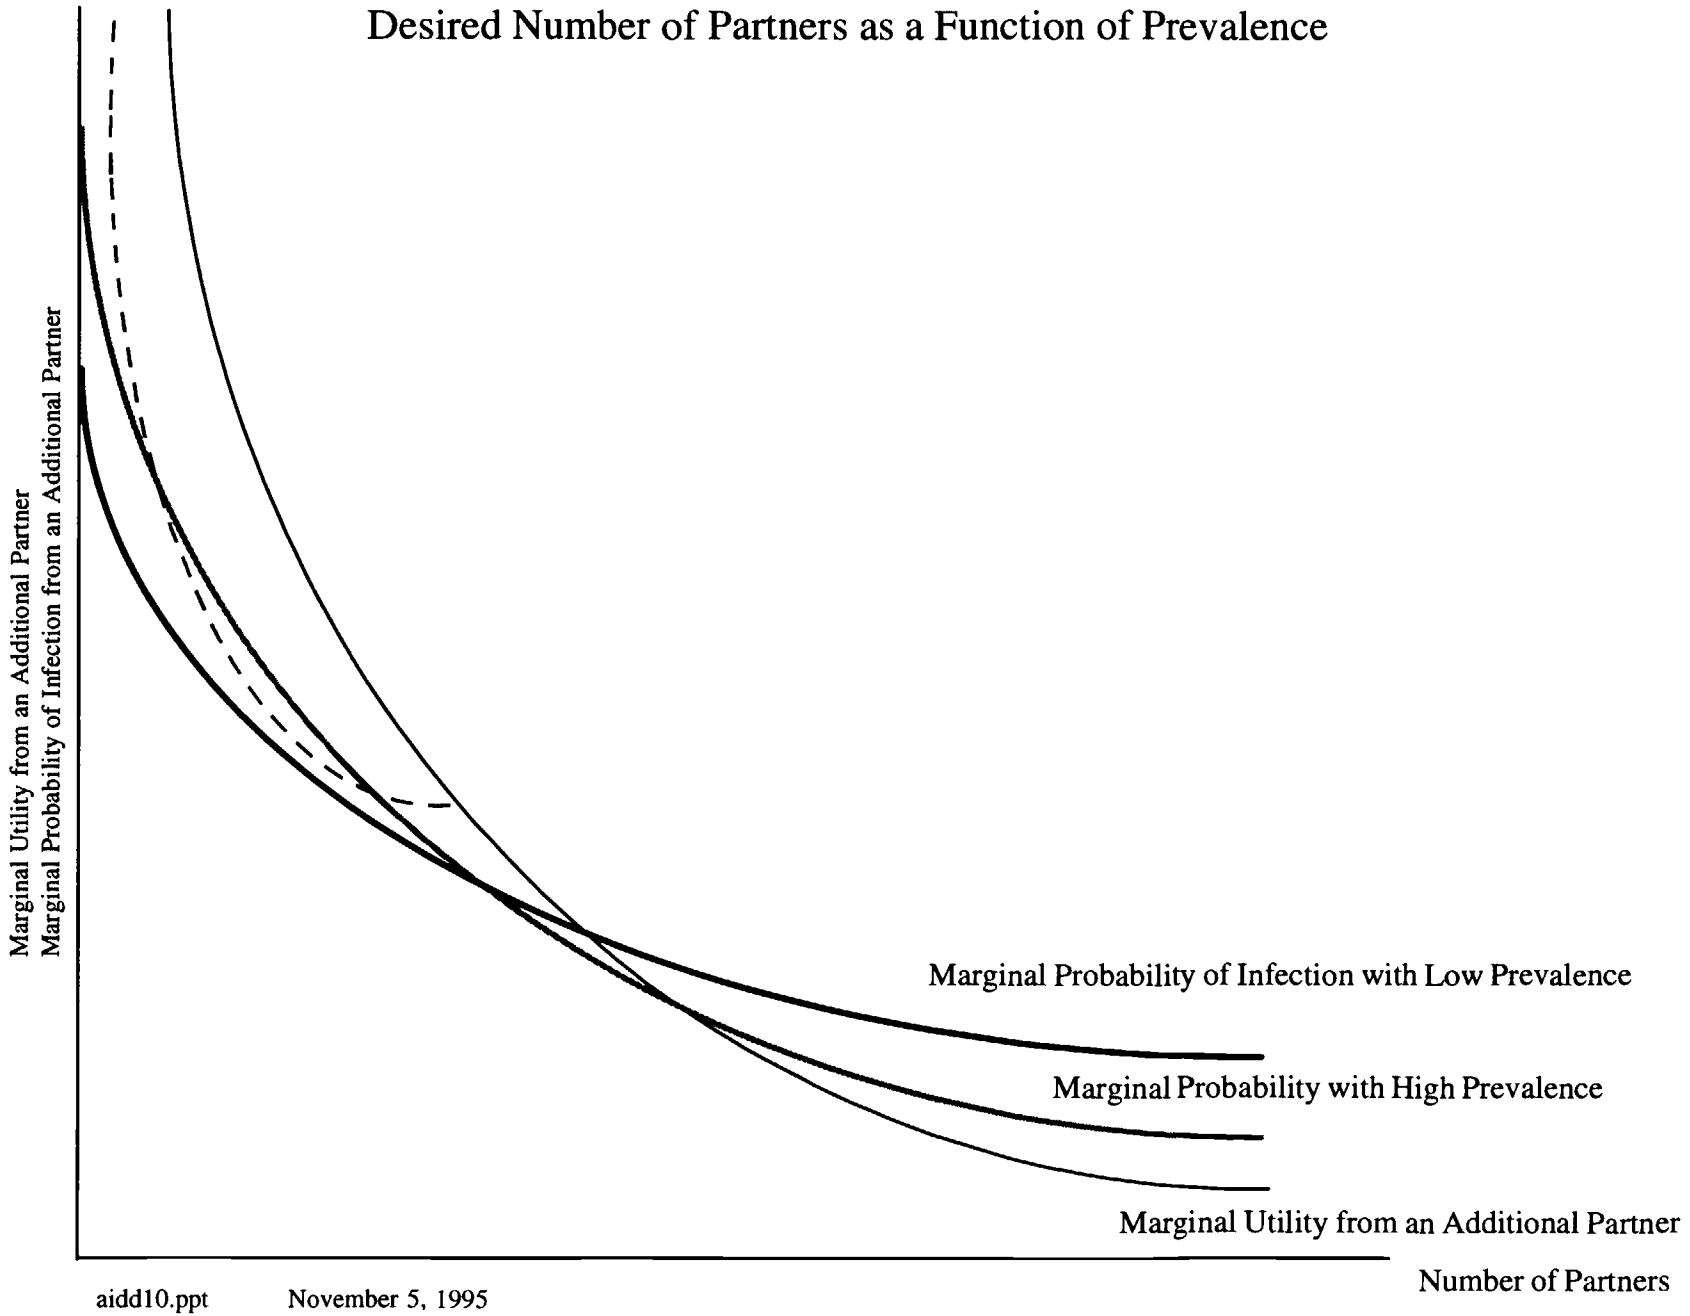

could go through even if some high-activity people reduced activity in response to increased prevalence, and others reduced activity to a low level.

II. Low-Prevalence, Heterogeneous Populations

The previous section considered the case of a population in which all people had the same i(Y) function. This section argues that if different people have different i(Y) functions, reductions in activity by low-activity people will raise weighted prevalence in the pool of remaining partners. If increased prevalence leads to greater proportional reductions in the rate of partner change by low-activity people than by high-activity people, weighted prevalence in the pool of partners will increase further, creating positive feedbacks and the possibility of multiple equilibria.[14]

Subsection II.A. identifies circumstances under which there will be multiple equilibria in instantaneous rates of partner change, given prevalence at a particular moment. Subsection II.B. shows that the interaction of epidemiological dynamics and behavioral choices may lead to multiple steady-state equilibria. In one steady-state equilibrium, low-activity people have few partners, high-activity people dominate the pool of available partners, and prevalence in the pool is therefore high. In the other steady-state equilibrium, low-activity people have a moderate number of partners and therefore constitute a majority of people in the pool of available partners, so that prevalence in the pool is low. Paradoxically, epidemiological

[14]The mechanism is similar to that in Akerlof's [1970] model of the market for lemons.

dynamics imply that prevalence may be higher in the steady-state equilibrium with lower activity.

A. Equilibria

Suppose that a proportion α_H of the population have i_H partners per period, and prevalence Y_H; and that the remaining proportion of the population, α_L has i_L partners per period and prevalence Y_L. If people mix randomly, the chance that a randomly selected partner will be infected is

$$W = \frac{\alpha_H i_H Y_H + \alpha_L i_L Y_L}{\alpha_H i_H + \alpha_L i_L}. \tag{3}$$

W, the percentage of couplings that involve an infected partner, is the probability of meeting a high-activity person, $\alpha_H i_H/(\alpha_H i_H+\alpha_L i_L)$, times the probability that a high-activity person will be infected, Y_H, plus the probability of meeting a low-activity person, times the probability that a low-activity person will be infected. I will refer to W as weighted prevalence in the pool of the available partners. W will in general be greater than Y, prevalence, because W weights people in proportion to their rate of partner change. Note that an increase in i_L, the rate of partner change of low-activity people, reduces W, at least in the short-run, because it increases the probability that a randomly selected partnership will be with a low-activity person.

This section examines the tractable case in which partners are selected randomly and

all people reduce their rate of partner change by equal absolute amounts in response to increased weighted prevalence in the pool. This would occur in a low prevalence population under the linear-quadratic utility function $u_k(i) = \theta_k i_k - \psi i_k^2$, where subscripts denote the individual's preference group. (I assume that a proportion α_L of the population has $\theta = \theta_L$, and a proportion $\alpha_H = 1 - \alpha_L$ has $\theta = \theta_H$.)[15] I will assume that it is possible to treat the probability of infection as linear in the number of partners, so that individuals maximize $u_k(i) - i\beta cW$, where c is the cost of infection. To simplify the algebra, I will normalize the cost of infection to $1/\beta$, so that individuals maximize $u_k(i) - iW$. Utility maximization then implies that

$$i_k = \max\left(\frac{\theta_k}{2\psi} - \frac{W}{2\psi}, 0\right). \tag{4}$$

Thus, if $i > 0$ for all groups, an increase of Δ in prevalence causes all groups to have $\Delta/2\psi$ fewer partners per period. This will increase weighted prevalence, W, as is clear from equation (3). It will thus spur further reductions in activity, further increasing W.

The assumptions of a linear-quadratic utility function and random matching are particularly tractable, and it is instructive to use them here at least to develop the basic concepts in the analysis. However, in Sections III and IV, I argue that similar positive

[15]Although I will refer to θ_H and θ_L as representing preferences for the number of partners, people may differ not in fundamental preferences, but in other variables which affect the costs and benefits of having additional partners. For example, if a prostitute needed more money because her child was sick, that would be referred to in this paper as an increase in θ, the parameter indexing her "preference" for the number of partners.

feedbacks can occur in high prevalence populations under an isoelastic utility function, in which people reduce activity equiproportionately in response to increases in the marginal probability of infection from an additional partner. Because the marginal probability of infection is less sensitive to prevalence for high activity people, these utility functions imply that low-activity people will make greater proportional reductions in activity in response to increased prevalence. More generally, under random matching, there will be positive feedbacks if low-activity people reduce their activity by a larger proportion than high-activity people in response to increased prevalence. Section IV argues that if high-activity people match with each other disproportionately rather than matching randomly, there can be positive feedbacks even if all people reduce activity equiproportionately in response to increased prevalence.

I will define a pair (i_L,i_H) as an *equilibrium*, given Y_H and Y_L, if weighted prevalence is W given activity levels i_H and i_L, and if weighted prevalence W induces low-activity people to have i_L partners per period and high-activity people to have i_H partners per period. A four-tuple (i_L,i_H,Y^*_L,Y^*_H) is defined as a *steady-state equilibrium* if (i_L,i_H) is an equilibrium number of partners given prevalence (Y^*_L,Y^*_H), and (Y^*_L,Y^*_H) is steady-state prevalence given activity (i_L,i_H). Equilibria are thus defined purely in behavioral terms, whereas steady-state equilibria are defined in both behavioral and epidemiological terms.

There can be multiple equilibria for moderate values of Y_H and Y_L, because the more low activity people enter the pool, the safer the pool will be, and therefore the more partners low-activity people will choose to have. However, if prevalence is sufficiently high, low-

activity people will never enter the pool and i_L will equal zero in the unique equilibrium.[16] Conversely, if prevalence is sufficiently small, i_L must be greater than zero in the unique equilibrium. Figure II shows the values of Y_H and Y_L for which multiple equilibria exist.[17] The results are summarized in the following proposition:

Proposition I.

1. If $Y_H < \theta_L$ there will be one equilibrium, in which $i_L > 0$.
2. If $[\theta_L(1+\alpha_H)- \alpha_H\theta_H- \alpha_L Y_L]/\alpha_H > Y_H > \theta_L$ and $[(\theta_L-\alpha_H(\theta_H-\theta_L)-(\alpha_L Y_L+\alpha_H Y_H)]^2 + 4\ \alpha_H(\theta_H-\theta_L)(\theta_L-Y_H) > 0$, then there will be one equilibrium with $i_L=0$ and two equilibria with $i_L>0$.
3. If $Y_H > \theta_L$ and $Y_H > [\theta_L(1+\alpha_H)- \alpha_H\theta_H- \alpha_L Y_L]/\alpha_H$, then there will be one equilibrium in which $i_L = 0$.

Proof of Proposition I[18]:

Note that if $Y_H< \theta_L$, there cannot be an equilibrium in which the low-activity people have no partners, because in that case weighted prevalence would be Y_H, which would induce low-activity people to have a positive number of partners.

In equilibria in which i_L (and thus i_H) are positive, the first-order conditions for choice of partners are satisfied with equality. Together with the definition of W, this implies that if both groups have a positive number of partners,

$$(5) \qquad i_L = \frac{\theta_L-\alpha_H(\theta_H-\theta_L)-(\alpha_L Y_L+\alpha_H Y_H)\pm\sqrt{(\theta_L-\alpha_H(\theta_H-\theta_L)-(\alpha_L Y_L+\alpha_H Y_H))^2+4\alpha_H(\theta_H-\theta_L)(\theta_L-Y_H)}}{4\psi}.$$

If $Y_H < \theta_L$, then the square root term is larger than the term outside the square root. Hence, one of the roots for i_L will be negative and there can be only one equilibrium, in which both groups have a positive number of partners.

On the other hand, if $Y_H> \theta_L$, then the square root term is smaller than the absolute value of the rest of the numerator. If $[(\theta_L-\alpha_H(\theta_H-\theta_L)-(\alpha_L Y_L+\alpha_H Y_H)]^2 + 4\ \alpha_H(\theta_H-\theta_L)(\theta_L-Y_H) > 0$,

---

[16]Note that if $\theta_L > 1$, so that low-activity people will have partners even if weighted prevalence equals one, there can be no $i_L = 0$ equilibrium.

[17]I assume that $Y_H > Y_L$, since the high group always has at least as many partners as the low group.

[18]I thank Cesaltina Pires for assistance with this proof.

Figure II

Region of Multiple Equilibria and Local Dynamics

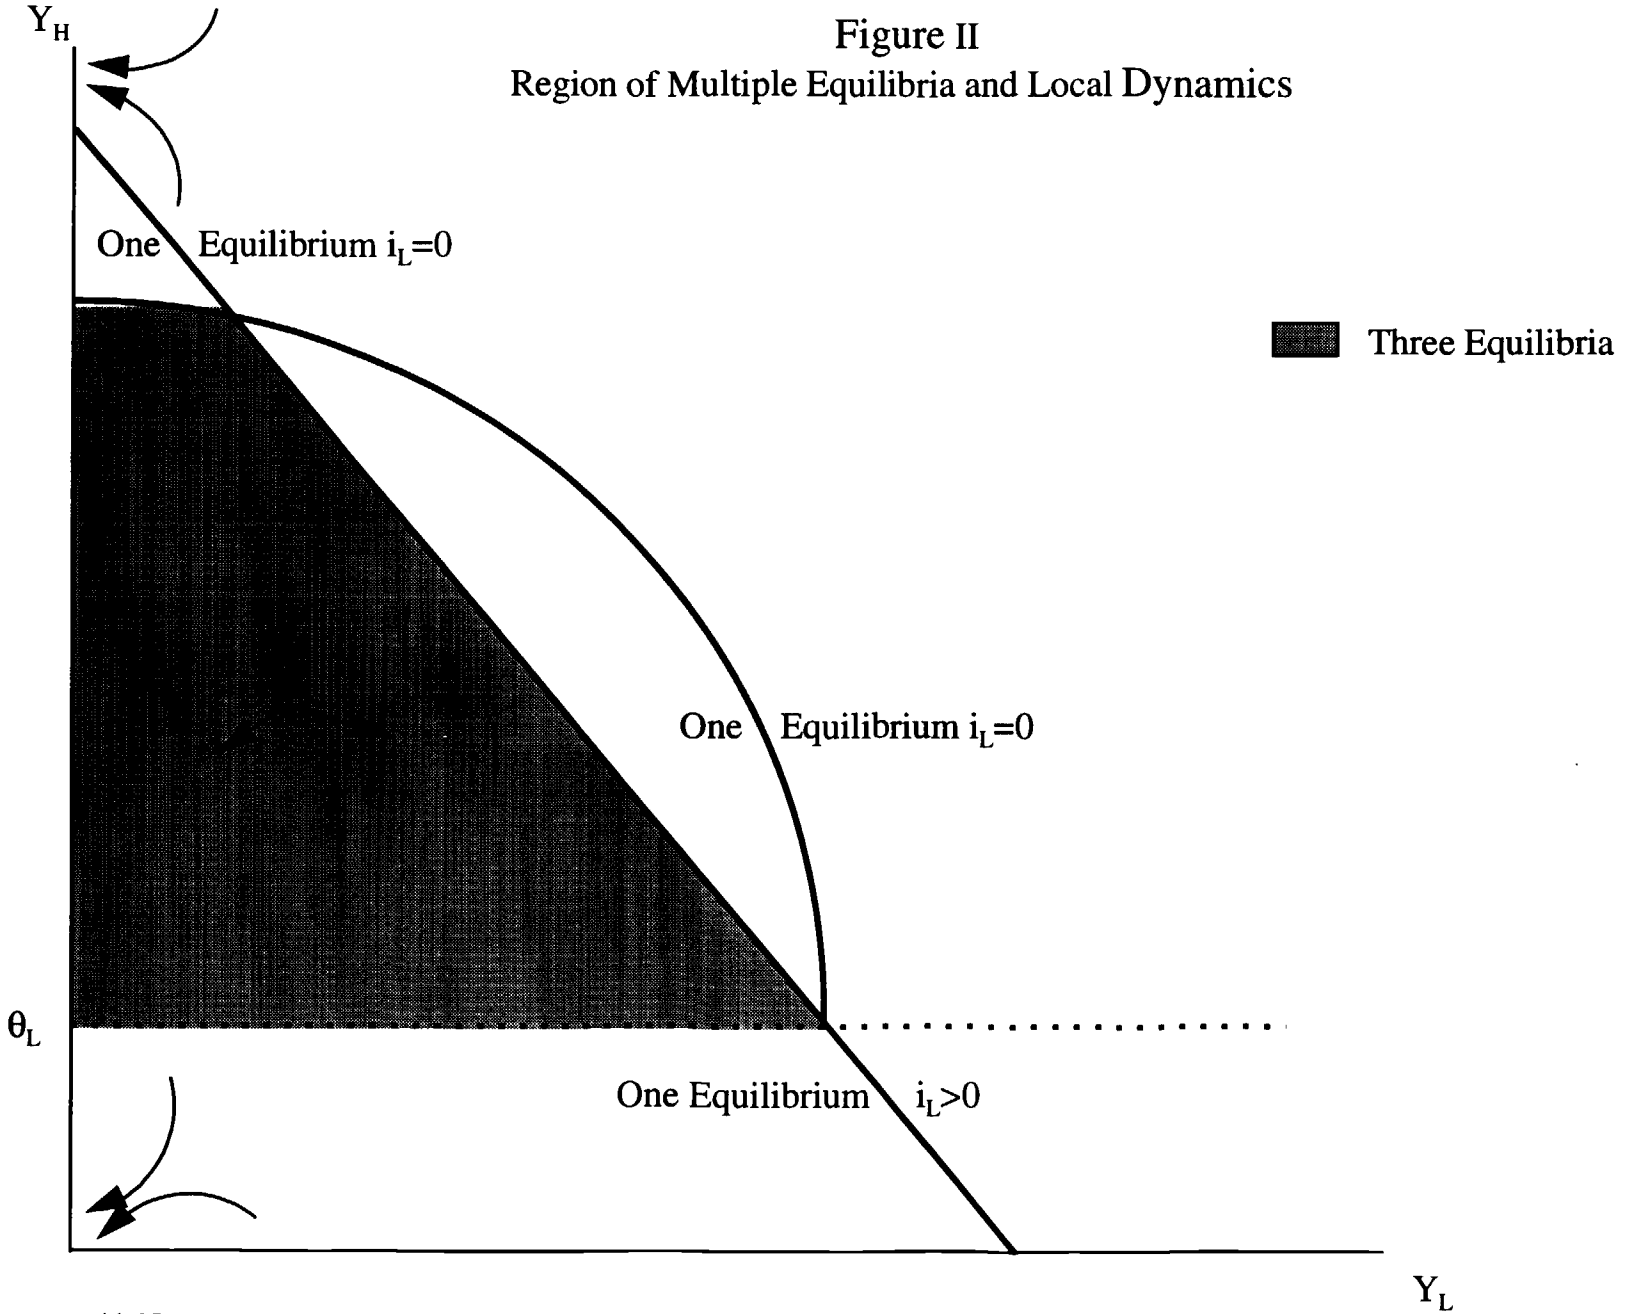

aidn05.ppt November 5, 1995

there will be two positive roots. If $Y_H > \theta_L$ and $Y_H > [\theta_L(1+\alpha_H) - \alpha_H\theta_H - \alpha_L Y_L]/\alpha_H$, the rest of the numerator is negative, and either there are two negative roots or two complex roots. Hence, in this case, there can be no equilibria with a positive number of partners in the low-activity group.

■

Equilibria with higher activity will be preferred by all people in the short run, since the higher the level of activity, the lower instantaneous weighted prevalence. The lower weighted prevalence, the better the tradeoff each individual faces between the number of partners and the risk of infection.[19]

B. Steady-State Equilibria

In this subsection, I argue that equal increases in activity by all people may reduce prevalence in the pool in the long-run, as well as the short-run. Thus, there may be multiple steady-state equilibria, as well as multiple equilibria. Moreover, steady-state equilibria with higher activity may have not only lower weighted prevalence, W, but lower unweighted prevalence, Y. In some cases, moving to an equilibrium with higher activity could lead to the eradication of the disease.

To see the intuition, consider an example in which an average of 10 partners per year were required for each infected person to infect one uninfected person before dying, and thus for the disease to survive. If the population consisted of a small minority which had 11 partners per year, and a majority which never had any partners, the disease would survive in the small group. Now suppose that the group with 11 partners per year increases to 11.02

[19]Higher activity equilibria will lead to more new infections, however, and I have not ruled out the possibility that they may lead to less preferred steady-state equilibria.

partners per year and that the group with zero partners per year increases to 0.02 partners per year, or approximately 1 partner over their lifetime. The high-activity people might match 6.01 times per year with other high-activity people and 5.01 times per year with low-activity people. High-activity people would infect low-activity people, but low-activity people would rarely pass the disease onto others. The high-activity people will only meet with each other 6.01 times per year instead of the 10 required for the disease to survive. Hence the disease would die out among the high-activity people, and thus among the population as a whole. To see this more formally,[20] recall that for the disease to spread, R_0, the basic reproductive rate of the disease, must be greater than one. In a randomly mixing population, in which a proportion α_k of the population has i_k partners per period and there are k activity groups,

$$\frac{dY_j}{dt} = i_j\beta(1 - Y_j)\sum_{k=1}^{K} \frac{i_k Y_k \alpha_k}{\sum_{k=1}^{K} i_k \alpha_k} - \delta Y_j. \tag{6}$$

Dividing by Y_j, taking the limit as prevalence approaches zero, and using the fact that at low prevalence Y_k / Y_j approaches i_k / i_j, Anderson and May [1991] show that in low-prevalence

[20]See Whittaker and Rentin [1992] and Kremer [1994] for a fuller explanation of why reductions in activity by low-activity people can lead to increased prevalence.

populations, R_0 will equal β/δ times a weighted average number of partners, c, where

$$c = \sum_{i=1}^{K} \left(i_K \frac{i_K \alpha_K}{\sum_{K=1}^{K} i_K \alpha_K} \right). \tag{7}$$

This weighted average number of partners can be expressed as $\mu + \sigma^2/\mu$, where μ and σ^2 denote the mean and variance of the number of partners per period.

Since $dc/d\mu = 1 - \sigma^2/\mu^2$, an equal increase in activity by all members of the population will reduce R_0 if $\mu < \sigma$ [Kremer, 1995]. Surveys of sexual activity suggest that this is typically the case. For example, a comprehensive study of sexual behavior in Britain, the National Survey of Sexual Attitudes and Lifestyles (NSSAL), found that the mean number of partners over the last five years was 1.98 and that the standard deviation was 4.36 [Johnson et al., 1994]. Simulations suggest that under the SI model, equal increases in activity by all members of the population might reduce R_0 below one and lead to the eradication of the disease in this population [Kremer, 1995].[21]

Since increases in activity may reduce prevalence, there may be one steady-state equilibrium in which activity is low, and weighted prevalence is high; and another steady-state equilibrium in which activity is high, and the disease is eradicated.

---

[21] Differentiating equation (7) with respect to i_k indicates that increases in activity by those with fewer than c/2 partners per period will reduce R_0, if activity of all other groups is held constant [Kremer, 1995]. This analysis suggests that increases in the rate of partner change by people with between c/2 and approximately c partners per period could potentially reduce steady-state prevalence by reducing weighted prevalence in the short-run, and thus encouraging people with fewer than c/2 partners per period to enter the pool.

Proposition II.

If

$$(8) \qquad \frac{(\alpha_L\theta_L^2 + \alpha_H\theta_H^2)}{(\alpha_L\theta_L + \alpha_H\theta_H)} < \frac{2\delta\psi}{\beta} .$$

then there will be a steady-state equilibrium with zero prevalence and high activity. If $2\delta\Psi/\beta < (\theta_H-\theta_L)(1-\theta_L)$, then there will be an equilibrium with low activity and high weighted prevalence. Hence, if $\alpha_L > \alpha_H$, there will be multiple steady-state equilibria for some values of $2\delta\Psi/\beta$, θ_H, and θ_L.

Proof of Proposition II.

There will be an equilibrium in which the disease is eradicated, $(i_H, i_L, Y_H, Y_L) = (\theta_L/2\Psi, \theta_H/2\Psi, 0, 0)$, if the first inequality in the proposition is satisfied, since

$$(9) \qquad R_0|_{\bar{Y}=0} = \frac{\beta}{\delta 2\psi} \frac{\alpha_L\theta_L^2 + \alpha_H\theta_H^2}{\alpha_H\theta_H + \alpha_L\theta_L} < 1 .$$

There will be an equilibrium with $i_L = 0$ if

$$(10) \qquad \frac{dY_H}{dt}|_{Y_H = \theta_L} \geq 0 ,$$

given that $i_L = 0$. At $Y_H = \theta_L$, $i_H = (\theta_H - \theta_L)/2\psi$, so (10) simplifies to the second inequality in the proposition above. It is straightforward to show that if $\alpha_L > \alpha_H$, there will be values of the other parameters such that both inequalities are satisfied.

A *robust steady-state equilibrium* (i_L, i_H, Y_L^*, Y_H^*) is defined as a steady-state equilibrium in which the equilibrium rate of partner change (i_L, i_H) is unique given prevalence (Y_L^*, Y_H^*). Thus once a society settles into a robust steady-state equilibrium, it will stay there. If $Y_H^* < \theta_L$, then, by Proposition I, there is only one equilibrium with i_L and $i_H > 0$, and the steady state is robust. Similarly, if there is a stable steady-state equilibrium in which the disease is eradicated, it will be robust. In contrast, the high-prevalence steady-state equilibrium in which low-activity people have no partners will be robust only if Y^*_H and α_H

are large enough to preclude any equilibrium with $i_L > 0$. For a large enough α_L, there will always be an equilibrium in which all of the low-activity people decide to enter the pool simultaneously, and weighted prevalence is therefore low enough to justify this decision. Figure III illustrates a situation in which there are two robust steady-state equilibria.

Proposition III.

The $i_L = 0$ steady-state equilibrium is robust if and only if

$$(\alpha_H\theta_H + \alpha_L\theta_L + Y_H^*\alpha_H)^2 - 4\theta_H Y_H^*\alpha_H < 0, \tag{11}$$

where Y^*_H is steady-state prevalence in a population consisting only of high-activity people, i.e. the solution to $\beta Y_H^{*2} - Y_H^*(\beta\theta_H + \beta) - \delta 2\psi + \beta\theta_H = 0$.

Corollary to Proposition III.

For any θ_L and θ_H, there will be a value of α_L, the proportion of the population in the low activity group, above which the steady-state equilibrium with $i_L = 0$ will not be robust.

Proof of Proposition III.

To check whether other equilibria are possible given that prevalence among the high-activity group is Y_H^*, note that if $Y_L = 0$, weighted prevalence is

$$W = \frac{i_h Y_H^*\alpha_H}{\alpha_H i_H + \alpha_L i_L}. \tag{12}$$

Substituting for i_H and i_L,

$$W = \frac{(\theta_H - W)Y_H^*\alpha_H}{\alpha_H(\theta_H - W) + \alpha_L(\theta_L - W)}. \tag{13}$$

This yields the quadratic equation, $W^2 - W(\alpha_H\theta_H + \alpha_L\theta_L + Y_H^*\alpha_H) + \theta_H Y_H^*\alpha_H = 0$. If $(\alpha_H\theta_H + \alpha_L\theta_L + Y_H^*\alpha_H)^2 - 4\theta_H Y_H^*\alpha_H < 0$ then the quadratic equation will have no real solutions, so the $i_L = 0$ equilibrium will be unique for $Y_L = 0$, $Y_H = Y_H^*$. Thus the $i_L = 0$ steady state equilibrium will be robust. If this expression is greater than zero, on the other hand, there

will be an admissible solution for W with $i_L > 0$, and the steady-state equilibrium will not be robust.

■

The possibility of multiple steady-state equilibria suggests that the timing of prevention efforts may be critical. For example, suppose that in the presence of prevention efforts, the set of equilibria is as depicted in Figure II, but that in the absence of prevention efforts, there would be no equilibrium in which the disease was eradicated. In this case, if the prevention efforts are started early enough, it may be possible to reach the equilibrium in which the disease is eradicated, but if health authorities wait until the disease is already established, it may be impossible to escape a steady-state equilibrium with high prevalence.

III. HIGH-PREVALENCE POPULATIONS

Previous sections considered the case in which β, the transmission rate, and i_{max}, the maximum rate of partner change, were low enough that all people would response to increased prevalence by reducing activity. This section considers the case in which increased weighted prevalence leads low-activity people to have fewer partners, but leads high-activity people to actually increase their number of partners. In this case, there may be one steady-state equilibrium in which high-activity people have very many partners, and low-activity people have very few partners, and another steady-state equilibrium in which the differences between the groups are smaller. High prevalence areas in which prostitution has continued, such as parts of Africa and Thailand, may be trapped in the former equilibrium.

I will assume that $u_k(i) = -\tau_k/i$ where τ_k is a preference parameter with value τ_L for a proportion of the population α_L and τ_H for a proportion $\alpha_H=1-\alpha_L$. This utility function implies that the elasticity of the desired number of partners to the probability of infection will be 1/2[22]. Changes in the marginal probability of infection prevalence will spur equiproportionate changes in activity by all groups. The desired number of partners for members of group k will be

$$i_k = \min\left[\frac{\delta}{\left(\sqrt{\frac{\delta}{\tau_k}} - \sqrt{\beta W}\right)\left(\sqrt{\beta W}\right)}, i_{max}\right] \tag{14}$$

for $\delta > \tau_k$ $\beta W > 0$, and i_{max} otherwise. Note that i_k declines in W until p(i)=1/2, and then increases to i_{max}. I will assume that i_{max} is sufficiently great that prevalence among people with i_{max} partners can be approximated as one. This may be roughly accurate for high-activity people in high-prevalence populations.

The epidemiological conditions for constant prevalence imply $\beta i_k(1 - Y_k)W = \delta Y_k$ for all groups k. A four-tuple (i_L, i_H, Y_L, Y_H) must satisfy both this epidemiological condition and the behavioral condition (14) to be a steady-state equilibrium.

---

[22]Increases in prevalence can only spur increased activity among high-activity people if the elasticity of the number of partners with respect to the marginal probability of infection from an additional partner is less than one. If the elasticity were equal to one, then an increase in prevalence would lead to an equiproportionate reduction in activity, leaving the lifetime probability of HIV infection constant. If the elasticity were greater than one, then increased prevalence in the pool of available partners would lead to reduced probability of infection. In fact, high-activity people might reduce their activity by a greater proportion than low-activity people in response to increased prevalence in the pool of available partners, because any initial reductions in activity would reduce their chance of becoming infected by an infra-marginal partner, and thus would further encourage them to reduce their activity.

The set of steady-state equilibria depend on τ_L and τ_H. For τ_L and $\tau_H > \delta/\beta$, $i_L = i_H = i_{max}$. For τ_L and $\tau_H << \beta/\delta$, both i_L and i_H will be less than i_{max}. The behavioral and epidemiological equations imply that any solution for steady-state prevalence in which i_L and i_H are less than i_{max} must satisfy[23]:

$$(15) \qquad W = \frac{\chi^2\psi^2}{2\beta\delta} - \alpha_H\sqrt{\frac{\tau_H}{\tau_L}} - \alpha_L\sqrt{\frac{\tau_L}{\tau_H}} \pm \frac{\chi\psi}{2\beta\delta}\sqrt{\chi^2\psi^2 - 4\beta\delta\left(\alpha_H\sqrt{\frac{\tau_H}{\tau_L}} + \alpha_L\sqrt{\frac{\tau_L}{\tau_H}}\right)},$$

where $\chi = \alpha_L\tau_L^{1/2} + \alpha_H\tau_H^{1/2}$ and $\psi = \beta + \delta(\tau_L\tau_H)^{-1/2}$.

For some parameter values such that $\tau_L << \delta/\beta < \tau_H$, it is possible to construct two steady-state equilibria. If $i_H(W) < i_{max}$, where W is given by equation (15), then there will be a steady-state equilibrium in which i_L and i_H are both less than i_{max}. There will also be a steady-state equilibrium in which $i_H = i_{max}$, so $Y_H \approx W \approx 1$, and

$$(16) \qquad i_L \approx i_L(1) = \frac{\delta}{\left(\sqrt{\frac{\delta}{\tau_L}} - \sqrt{\beta}\right)\sqrt{\beta}},$$

if high-activity people dominate the pool if they have i_{max} partners, i.e., if $\alpha_H i_{max} >> \alpha_L i_L(1)$. Thus there may be one steady-state equilibrium with low variance in activity and another with high variance.

Social welfare is greater in the low-variance steady state. This is because weighted

[23]Calculations, for which I thank Sergei Severinov, are available from the author.

prevalence is lower in the low-variance steady state, and people will therefore face a better trade-off between the probability of infection and the desire for a certain number of partners. Nonetheless, the high-variance steady-state could have lower unweighted prevalence, especially if the high-activity group is a small minority. For example, in the high-variance steady-state equilibrium, a tiny minority of the population could have i_{max} partners and the vast majority could have very few, so that unweighted prevalence in the population as a whole could be low.

The model may help explain why prostitution has continued in high-prevalence areas, because it suggests that if prevalence is already high, further increases in prevalence among men may reduce activity among low-activity women, but spur prostitution.[24] In a one-sex model, increases in activity by high-activity people directly increase the weighted prevalence that all other people face in the pool of available partners. In a two-sex model, the causal chain is more complicated. Increases in activity by high-activity members of one sex directly worsen the pool of available partners for the opposite sex. This increases prevalence among the opposite sex, and thus worsens the pool for other members of their own sex.

Consider an extreme case in which all men have the same level of activity; women are divided into prostitutes and non-prostitutes; and prostitutes have a 100 percent chance of infection. If prevalence increases among men, sex becomes more dangerous for women who are not prostitutes and they therefore reduce their sexual activity. However, the increase in

---

[24]Anderson and May [1991] argue that if public health campaigns reduce sexual activity among low-activity women, substitution by men towards high-activity women could cause an increase in prostitution. This section argues that increases in prevalence among men may create incentives for exactly this behavior.

prevalence among men does not make sex more dangerous for prostitutes, since they already had a 100 percent chance of being infected. Similarly, sex with prostitutes does not become more hazardous for men, and in fact the increase in prevalence among men lowers their marginal risk of further sexual activity. Hence if men substitute at all between prostitutes and non-prostitutes, increases in prevalence among men will lead to an increase in the demand for prostitutes.

In this extreme example, in which becoming a prostitute entails a 100 percent chance of infection, an increase in prevalence among men also creates incentives for an increase in the supply of prostitutes. This is because the alternative to being a prostitute is not a zero chance of infection, but the possibility of infection by a partner who has had many previous partners, or is not monogamous. This possibility increases with prevalence among men.[25] Since increases in infection rates among men increase both demand for prostitution among men, and supply among women, they will spur prostitution, causing further increases in prevalence, leading to a vicious cycle.

At less than 100 percent prevalence among prostitutes, the results depend on how elastic prostitute demand and supply are to the probability of infection, price, and the behavior of women who are not prostitutes. Even if the parameters are not such that increases in prevalence among men actually spur prostitution, they may well be such that increases in prevalence have only a limited impact on prostitution in situations in which many women's other alternatives are extraordinarily bleak.

---

[25]Moreover, as Philipson and Posner point out, some proportion of women who know they are infected may become prostitutes.

IV. PREFERRED MATCHING

In the previous sections, I assumed that people chose sexual partners randomly. However, if people have some information about potential partners' sexual history, high-activity people may be disproportionately likely to match with each other [Kaplan, Cramton, and Paltiel, 1990]. This can be modeled using the preferred matching approach of Jacquez et al. [1988], in which people match within their group with probability γ and match randomly with probability $1-\gamma$.

There is greater scope for positive feedbacks under preferred matching than under random matching. Recall that under random matching, there could only be positive feedbacks in low prevalence populations if low-activity people reduced their activity by a greater proportion than high-activity people in response to increased prevalence. Under preferred matching, changes in activity will create positive feedbacks even if all groups respond to a given percentage increase in the marginal probability of infection with an equal percentage reduction in activity.

To see this, note that under preferred matching, the probability that the partner of someone in the jth group will be infected is $\gamma Y_j + (1-\gamma)W$, where Y_j denotes prevalence in the ith group and W denotes weighted prevalence. For low-activity people, the main risk of infection comes from meeting people randomly in the pool, whereas for high-activity people, the main risk of infection come from matching within their own group.

Suppose that low-activity people increase their activity so that prevalence in the pool

falls. The marginal probability of infection from an additional partner changes by a small proportion for high-activity people, and hence they will only increase their rate of partner change by a small proportion. However, for the low-activity people, the marginal probability of infection from an additional partner falls by a large proportion. Hence low-activity people will increase their rate of partner change. Prevalence will thus fall further, setting off another round of indirect effects.

However, in low-prevalence populations, these positive feedbacks will not be strong enough to create multiple equilibria if the elasticity of the number of partners to the marginal probability of infection from an additional partner is less than one.[26] To see this, denote the preferred number of partners for a low-activity person as $i_L^*(\bar{i}_L, i_h)$, where $\bar{i}_L$ denotes the number of partners of all other low-activity people. Since there will be a unique preferred level of i_H for any $\bar{i}_L$, this can be rewritten as $i_L^*(i_H(\bar{i}_L), \bar{i}_L)$ or $i_L^*(\bar{i}_L)$. Since for isoelastic utility and $Y_L > 0$, this function is continuous, bounded, and has a positive intercept, $di_l^*/d\bar{i}_L$ must be greater than one at some point for there to be multiple equilibria.[27] An x percent increase in $\bar{i}_L$ will necessarily reduce prevalence in the pool by less than x percent. Hence, if the elasticity of demand for partners with respect to the risk of infection is less than or equal

---

[26]I have not proven that there can be multiple equilibria if the elasticity is greater than one, and I have not established necessary and sufficient conditions for multiple steady-state equilibria under preferred mixing.

[27]This is a clear implication of Cooper and John [1988] although they do not explicitly note this.

to one, this x percent reduction in prevalence will spur less than an x percent increase in $\overline{i_L}$. Thus, there cannot be multiple equilibria in low-prevalence populations if the elasticity is less than or equal to one.[28]

In high-prevalence populations, high-activity people are likely to be less sensitive to changes in the probabilities of infection than low-activity people. This will increase the scope for multiple equilibria, as argued in Section III.

V. CONCLUSION

Previous authors have argued that behavioral response to AIDS will counteract public policies to fight the epidemic [Philipson and Posner, 1993; Hadeler and Castillo-Chavez]. In contrast, this analysis, which models people making decisions about their rate of partner change given the composition of the pool of available partners, suggests that high-prevalence populations, behavioral response may increase the effectiveness of policy.

Whereas increased prevalence will lead low-activity people to reduce their sexual activity, it may lead high-activity people to become fatalistic, and thus either reduce their activity only slightly or actually increase it. In either case, the composition of the pool of available partners will worsen, creating a positive feedback. If this effect is strong enough, there may be multiple equilibria in sexual activity for given prevalence.

The model implies that public health messages will have a bigger effect if they are targeted to prevent high-activity people from exceeding the threshold level of activity at

---

[28]A more formal proof, by Robert Shimer, is available from the author.

which they are so likely to be infected that they choose to have even more partners. This will involve targeting highly active people, although perhaps not the most active, whose behavior may be very difficult to change. Reductions in transmission rates among people below the threshold level of activity will be partially counteracted by behavioral changes. On the other hand, above the threshold level of activity, reductions in transmission rates induced by public health measures will lead to changes in behavior that further reduce prevalence. The model also suggests that it may be exaggerating the risk of infection in high-prevalence groups may spread fatalism.

The model implies that below some threshold level of prevalence, society may approach a steady state in which differences in sexual activity among the population are modest. Above the threshold, prevalence might approach a steady state in which some members of the population have a very high rate of partner change, and others have a very low rate. Although prevalence is not necessarily lower in the high-variance steady state, welfare is lower since each individual faces a worse tradeoff between the probability of infection and the desired number of partners. Early public health campaigns that keep society from crossing the threshold level of prevalence could allow society to avoid this steady-state.

References

Ahituv, Avner, V. Joseph Hotz, and Thomas Philipson, "Will the AIDS Epidemic be Self-Limiting? Evidence on the Responsiveness of the Demand for Condoms to the Prevalence of AIDS," working paper, University of Chicago, 1993.

Akerlof, G. "The Market for Lemons: Qualitative Uncertainty and the Market Mechanism," Quarterly Journal of Economics, LXXXIV(1970): 488-500.

Anderson, Roy M. and Robert M. May, Infectious Diseases of Humans: Dynamics and Control, (Oxford: Oxford University Press, 1991).

Anderson, Roy M., Robert M. May, M.C. Boilly, G.P. Garnett, and J.T. Rowley, "The Spread of HIV-1 in Africa: Sexual Contact Patterns and the Predicted Demographic Impact of AIDS," Nature, (August 1991), 352-415.

Avery, Chris, Jody Heymann, and Richard Zeckhauser, "Individual Choice and the Spread of Infectious Diseases," working paper, Kennedy School of Government, 1994.

Barnett, Tony and Piers Blaikie, AIDS in Africa: Its Present and Future Impact, (New York: Guilford Press, 1992).

Becker, Gary, Treatise on the Family, (Cambridge, MA: Harvard University Press, 1981).

Bloom, David E. and Ajay S. Mahal, "Does the AIDS Epidemic Really Threaten Economic Growth?" National Bureau of Economic Research, working paper no. 5148, 1995.

Boozer, Michael and Tomas Philipson, "The Private Demand for Information and the Effects of Public Testing Programs: The Case of HIV," working paper, Yale University, 1994.

Castillo-Chavez, Carlos and Hadeler, K.P., "A Core Group Model for Disease Transmission," Working Paper, Cornell University, 1994.

Cooper, Russell and Andrew John, "Coordinating Coordination Failures in Keynesian Models," Quarterly Journal of Economics, CIII (August 1988), 441-64.

Geoffard, Pierre-Yves and Tomas Philipson, "Disease Eradication: Private vs. Public Vaccination," forthcoming, American Economic Review, 1995a.

Geoffard, Pierre-Yves and Tomas Philipson, "The Empirical Content of Canonical Models of Infectious Diseases: the Proportional Hazard Specification," Biometrika, 82 (1), 1995b 101-111.

Geoffard, Pierre-Yves and Tomas Philipson, "Rational Epidemics and Their Public Control," forthcoming, International Economic Review, 1995c.

Green, Jesse, "Who Put The Lid on gp120?" New York Times Magazine, (26 March 1995), 51-57.

Hethcote, Herbert W. and J.A.Yorke, "Gonorrhea: Transmission Dynamics and Control," Lecture Notes in Biomathematics, 56(1984), 1-105.

Jacquez, J.A., C.P. Simon, J. Koopman, L. Sattenspiel, and T. Perry, "Modelling and Analyzing HIV Transmission: the Effect of Contact Patterns," Mathematical Bioscience, 92(1988), 119-199.

Johnson, A.M., J. Wadsworth, K. Wellings, and J. Field, Sexual Attitudes and Lifestyles, (Oxford: Blackwell Scientific, 1993).

Kaplan, E.H., P.C. Cramton, and A.D. Paltiel, "Nonrandom Mixing Models of HIV Transmission," Mathematical and Statistical Approaches to AIDS Epidemiology,Lecture Notes in Biomathematics, C. Castillo-Chavez, ed., (Berlin: Springer-Verlag, 1990).

Kremer, Michael,"Can Having Fewer Partners Increase Prevalence of AIDS?" National Bureau of Economic Research Working Paper No. 4942, 1994.

Kremer, Michael, "AIDS and Imperfect Signals of Risk", MIT, unpublished notes, 1995.

May, Robert M. and Roy M. Anderson, "Transmission Dynamics of HIV Infection," Nature, 326(1987),137-142.

McKusick, L., Wiley, J. A., Coates, T.J., and Stall, G. (1985). Reported changes in the sexual behavior in men at risk from AIDS, San Francisco 1982-84: the AIDS Behavioral Research Project. Public Health Rep., 100(1985), 622-9.

Over, Mead and Peter Piot, "HIV Infection and Sexually Transmitted Diseases," Disease Control Priorities in Developing Countries, Dean T. Jamison and W. Henry Mosley, eds., (New York: Oxford University Press, 1993).

Philipson, Tomas, "Private Vaccination and Public Health: A Theoretical and Empirical Examination,"forthcoming, Journal of Human Resources, (1995). XXX

Philipson, Tomas and Richard Posner, Private Choices and Public Health: The AIDS Epidemic in an Economic Perspective, (Cambridge,MA: Harvard University Press, 1993).

Piot, P., M. Laga, R. Ryder, J. Perriens, M. Temmerman, W. Heward, and J.W. Curran., "The Global Epidemiology of HIV Infection: Continuity, Heterogeneity, and Change," Journal of AIDS, 3(1990), 403-11.

Velasco-Hernandez, Jorge X., Fred Brauer and Carlos Castillo-Chavez, "Effects of Treatment and Behavioral Change on the Dynamics of HIV/AIDS in a Highly Homosexually-active Population," University of Wisconsin, unpublished, 31 May 1994.

Velasco-Hernandez, Jorge X., Fred Brauer and Carlos Castillo-Chavez, "Recruitment Effects in Heterosexually Transmitted Disease Models," Cornell University, unpublished, August 1993.

Wellings, Kaye, Julia Field, Anne M. Johnson, and Jane Wadsworth, Sexual Behavior in Britain, The National Survey of Sexual Attitudes and Lifestyles, (London, England: Penguin Books, 1993).

Whitaker, L. and A.M. Rentin, "A Theoretical Problem of Interpreting the Recently Reported Increase of Homosexual Gonorrhea," European Journal of Epidemiology, (March 1992),187-191.

WHO (World Health Organization), The Current Global Situation of the HIV/AIDS Epidemic, (Geneva: WHO, Global Programme on AIDS, 1995).