Marital Shopping and Epidemic AIDS

Jeremy R. Magruder*

June 19, 2008

Abstract

Women over the age of 30 are at greatly reduced risks of HIV in South Africa, despite no reported declines in coital frequencies or pregnancies at this young age. Many existing models for the spread of HIV suggest a risk profile which is either independent of or increasing in age, and I calibrate a model of age-independent heterogeneous risk to show that only a model with age-decreasing risks can generate the observed age-death profile. Adapting Jovanovic's (1979) job-matching model into a model of spousal search, I find that serially monogamous spousal search is capable of generating epidemic prevalence levels despite very low average transmission rates, because search behavior interacts with peculiarities in the dynamics of HIV infectiousness. The implied age-infection profile closely mimics that observed in South Africa for both men and women in a way insensitive to parametric assumptions. Spousal search is also found to be consistent with micro-empirical evidence on HIV across Africa. Finally, I evaluate two policies: condom usage in short relationships and reducing transmission rates through biological means. Both are found to be somewhat effective, and used together almost completely eliminate the potential of spousal search to spread HIV.

1 Introduction

A well-established consensus within the medical literature indicates that the average transmission rate of HIV is extremely low, on the order of 1 in 1000 per sex act (e.g. Gray et al 2001, Quinn et al 2000, Fideli et all 2001), or 8-12% per partner-year. Despite this, antenatal prevalences reach as high as 40% in Botswana and 25% in South Africa, suggesting that a large fraction of the population have been afflicted by this extremely unlikely event.

*University of California, Berkeley. e-mail jmagruder@berkeley.edu. I am indebted to Fabian Lange, Ethan Ligon, Ted Miguel, Nicoli Nattrass, Emily Oster, Rohini Pande, Paul Schultz, Jeremy Seekings, Chris Udry, Damien de Walque, and seminar participants at Yale, Berkeley, and the University of Cape Town for many helpful suggestions. All mistakes are naturally my own.

Unsurprisingly, social scientists, epidemiologists, and the popular press have each weighed in on what sorts of behavior could yield such an outcome, with behavioral models ranging from heterogeneous preferences over risk (Kremer 1996) to large networks of concurrent sexual partners (Morris and Kretzschmar 1997) to unfaithful husbands (Kristof 2005).

This paper asks how most people become infected with HIV, and proposes that the age-profile of infection is a useful aggregate statistic to assess whether any of these models could predict most infections. Examining death registries in South Africa, a strong age-trend emerges – HIV infection risk appears to decline very sharply after the age 30 or so for women (and 35 for men). I open this paper by discussing the age-profile of infection which we can infer from the South African death registries, and illustrate that this decline in risk occurs too early in life to be explicable simply by age-changing preferences, at least if these preferences are correlated with reported coital frequencies or pregnancy rates. I then argue that we can parse most existing social science HIV models into those which predict increasing risks with age and those which suggest age-independence (though a spurious age-decline may be generated by the absorbing nature of HIV infection, particularly in the latter case). Intuitively, the sharp decline in risk levels with age seems unlikely to be generated by behavior where risk is age-independent or age-increasing; as a formal test, I calibrate a model of age-independent heterogeneous risk and find that this class of models do not fit the age-death profile in South Africa well, suggesting that heterogeneous risk models do not explain how most people behave[1]. As these models dominate age-increasing risk in terms of fitting the observed decline in HIV infections with age, rejecting the age-independent model rules out this latter class of models as well.

Instead, the age-death profile prefers a model of decreasing risk, with risks of new infection declining around age 30 for women (shortly after the median age of first marriage). This paper is the first to show that a simple search and matching model of serial monogamy is

[1]Elements of these models may all the same have severe epidemiological consequences, as variations in the behavior of a small fraction of individuals may cause huge swings in prevalence. I return to this point in the conclusion.

capable of generating high prevalences with relatively few partners despite a low average transmission rate due to peculiarities in the dynamics of HIV transmissivity. That is, epidemic AIDS is possible even when the vast majority of individuals do not engage in high risk behaviors. Using DHS data on marital behavior and theoretical predictions of search-and-matching models, I show that spousal search generates an age-death profile which fits that observed in South Africa well in a way robust to a large variety of parametric assumptions. I illustrate that DHS data from across Africa support several predictions of matching models, including the relationship of HIV with marital tenure and length of singlehood. Given that reported behavior looks similar in the US and Cape Town, I show that the variance in annual sexual partners is more consistent with matching than static preferences and examine whether the marital shopping model can predict the age-profile of Gonorrhea in the United States. Finally, I evaluate the public health implications of Marital Shopping as an important vector for the spread of HIV by examining policies of transmission rate reductions and short-term condom usage.

2 HIV in South Africa

Much of the data which could be used to learn about covariates of HIV infection is suspect. Ante-natal prevalence data is subject to selection bias, population survey data tend to have high refusal rates, and cause of death statistics may be rigged. Changes in overall death rates in areas where the epidemic hit suddenly and recently, however, are immune to these concerns. Figures 1 and 2 illustrate the deaths by age and gender observed in South Africa, where death registration is relatively complete. Epidemic AIDS is a recent phenomenon in South Africa, where ante-natal prevalences were 0.7% as recently as 1990 (RSA 2008). This is reflected in the bottom line of each figure, which shows a relatively flat deaths-by-age profile in 1996. The top line, in contrast, tells a different story. By 2002, the number of 30-40 year old women who died had more than tripled in six years. Looking out into

older ages, we see no such difference. There is a much slighter increase in deaths for 40-45 year old women, and by age 50, the death rate looks very similar to the historical one. For men, the age-death profile also peaks, though this happens five years later and the decline is less sharp. If we knew the precise time path of prevalence in South Africa, we could infer at precisely which age infection risk peaks. We don't; however, a close approximation can be found simply by subtracting 10 years from the death rates as individuals survive, on average, about ten years after infection. This approximation suggests that women are becoming much safer after the age of 30[2].

As this paper will focus on behavioral models which can predict such an age-profile, it is important that this is not an immediate outcome of age-changing preferences over coital frequencies. Sexual behavior data is notoriously difficult to measure, and individuals, particularly women (Gersovitz et al 1998) have been shown to misreport badly. Still, thirty seems a very young age for a sharp decrease in preferences over coital frequency, and two forms of evidence are available to shed some light on this issue, both available in the South African DHS. The first is self-reports of sexual behavior, which is presented in Table 1. For all potential coital frequencies, we see reported sexual behavior increasing until age 40, after which it remains higher than the behavior of women in their 20s, which are the ages that the death profile tells us involve peak HIV risk.

Given the evidence on misreporting, we may be concerned that these differences simply reflect differences in reporting bias. As such pregnancy is preferred as a biomarker. Table 2 reports pregnancy rates by age. Indeed, pregnancy rates do decline with age after age 30 as reported in column 1. However, there is a vast literature documenting a natural decline in fecundity beginning at age 30 or so (reviewed in te Velde and Pearson 2002). Table 2 shows

[2]This trend is not unique to South Africa, nor is it a recent observation. Though good mortality records do not exist for other countries in sub-Saharan Africa, authors in individual medical studies have found similar age-death profiles from HIV in Uganda (Sewankambo et al 2000) and the Democratic Republic of Congo (Pictet et al 1998). Anderson et al (1991) provided an early documentation of this age-pattern in an influential review. Despite the long-standing awareness of this trend, previous work has not used it to assess the plausibility of behavioral models to the author's knowledge.

births per year divided by the probability of conception for that age found in two studies, Templeton et al (1996) and van Noord-Zaadstra et al (1991)[3]. As table 2 shows, the actual births per year estimates seem to be in between what the two estimates would predict for women if sexual behavior remained constant, at least for women under 40. Sample sizes are tiny in both studies for women over 40, so these may be unreliable. While the above table cannot speak to individual tastes for variety, if these tastes are correlated with tastes for frequency, then we can say with some confidence that they are not declining with age, at least through age 40. These data do not support the hypothesis that women's age-specific preferences over coital frequency are insulating them from their husbands' behavior at older ages, nor that women begin to prefer lower coital frequencies at an early age. As such, I will follow the reports of behavior and pregnancies and assume in each model that the death profile from ages 10-50 (primarily reflecting behavior before age 40) is not determined by age-changing preferences over coital frequency.

3 Non-behavioral models of HIV

For the age-trend in infection observed in South Africa to be a suitable criteria to evaluate behavioral models, it is important that it is not simply determined by a biological component. Other sexually transmitted infections increase the transmission risk of HIV (e.g. Oster 2005) which may have age-implications. In fact, the observation that young people are far more likely to contract sexually transmitted disease is not unique to HIV or to Africa. Whether Gonorrhea or Chlamydia in the US (CDC 2005) , Human Papilloma Virus in Costa Rica (Castle et al 2005), or Herpes Simplex Virus-2 throughout the world (Smith and Robinson 2002), the young are consistently the population who contract STDs. Despite this observation, a closer reading of the medical literature reveals that the presence

[3]Templeton considers all in-vitro fertilizations in Britain from 1991-1994, while van Noord-Zaadstra considers all artificial inseminations in two fertility clinics in the Netherlands during the clinic-specific period when fresh (rather than frozen) semen was used (1973-1980 for one clinic, 1973-1986 for the other), and restricts analysis to married women whose husbands were azoospermic and who had never previously given birth or received an artificial insemination.

of other sexually transmitted diseases actually suggests an age-increasing risk profile for HIV infection. Though there is mixed evidence on the role of non-ulcerative STDs for HIV transmission, it is widely believed that ulcerative sexually transmitted diseases play a much larger role (e.g. Fleming and Wasserheit 1999) (and non-ulcerative diseases may have no effect at all, see Quinn et al 2000 and Gray et al 2001). The most prevalent ulcerative STD in Africa and the world is Herpes Simplex Virus -2 (e.g. Chen et al 2000, Wawer et al 1999). Therefore, if STD risk alone is explaining the age-profile of HIV in Africa, we can examine the age-profile of HSV-2 to determine the suggested age-risk pattern for HIV. Figures 3 and 4 (Smith and Robinson 2002) reveal that prevalence is highest among older people (as one might expect for a disease which is incurable and non-fatal), so from STDs alone older men and women should be at an elevated risk of HIV contraction, not a much lower one. STDs do not serve as a stand alone story for HIV in Africa; rather, they increase the need for an accompanying behavioral story to both explain the age-pattern of their spread and to further explain why those at the ages of greatest risk from STD prevalence are not contracting HIV.

4 Existing behavioral models of HIV

Because average HIV infectiousness is so low, models which would strive to achieve high HIV prevalences with feasible amounts of sexual behavior are often forced to rely on fundamental heterogeneity between individuals. The class of models which has dominated the economics literature revolves around individuals having heterogeneous preferences of sexual risk or variety (e.g. Kremer 1996, Philipson and Posner 1993). In Kremer's specification, heterogeneous individuals optimize a rate a partner change. A high risk group prefers to seek out many partners, while lower risk individuals only occasionally search for new partners. If there is a random component to matching, then high risk individuals are overrepresented in the pool of available partners due to their greater frequency of partner search, increasing the risk of drawing a dangerous partner above the population-average level.

This model belongs to the class of age-independent risk models. While different individuals prefer different risk levels and may be probabilistically likely to become infected earlier in life than others, an uninfected individual will face an identical probability of infection each period in which he remains uninfected. Another model which conforms to this class is the model of sexual networks laid out by Morris and Kretzschmar (1997). Morris and Kretzschmar argue that in order to achieve high prevalences, individuals must belong to large sexual networks of concurrent partners. Once HIV is introduced to a network, it can spread quickly throughout the network, exposing all members to high risk. As such, all members of a network are exposed to very similar risk levels, regardless of age. Presumably, individuals are heterogeneous with respect to the size and turnover rate of their networks, so we may expect this model to produce very similar age predictions to the preference-based model[4].

Can an age-independent risk model use heterogeneous types to generate an age-death profile like the one observed in South Africa? In principle, a two type model might be very good at it: high risk types become quickly infected, and die out while still young. Low-risk types, in contrast, remain safe throughout. However, this sort of explanation makes two extreme assumptions: first, there are no "middle risk types," who continue to be probabilistically infected later in life; and second, that the epidemic is in its steady state. Self-reports of the number of sexual partners in the last year from 14-22 year olds in South Africa, though subject to all of the concerns discussed in the previous section, seems to contradict the first assumption[5]. As reported in columns 2 and 4 of table 3, young adults report a triangular distribution of annual partners, with the majority being in a low-risk category. Intuitively, if heterogeneous risk is distributed so that a large number

[4]One may imagine many ways by which network formation may covary with age. However, the assumptions needed for networks to create age-decreasing risks are strong. All members of a network are at similar levels of risk, so if network members are both young and old then we have age independence. If networks of partners are stratified by age, one could imagine a situation where uninfected older individuals match separately from more infected younger people. However, the dynamics of these networks are very difficult to justify. For example, if everyone in one network matches only with people born before January 1, 1970 then such a network could exist, but not if everyone matches with people who are within one or two years of their age.

[5]The South African Data Source for this table is wave 1 of the Cape Area Panel Study, which is discussed in the data appendix

of infections happen to low-risk individuals, then these individuals are still being infected later in life and would generate a different age-death profile. More importantly, epidemic stability is necessary as this age-independent explanation revolves around the high-risk older women having already attrited from the population, which would not have happened at the onset of the epidemic (for example, in South Africa's case).

4.1 Age-Independent Model Calibration

To evaluate whether an age-independent model is capable of generating the age-death profile observed in South Africa, I calibrate a model with up to three types. Each type is indexed by an annual risk level so that one is at high risk, one at low risk, and one at negligeable risk. It can be shown in a similar calibration (results available from the author) that if the majority of at-risk people are in the low-risk type, or the high risk type is not extremely high risk, only a bad fit can be generated in the steady-state. Here, I illustrate that regardless of the distribution of these three types, the epidemic is too young in South Africa to generate an age-death profile like the one observed. To do so, let $\rho_g(y)$ be the infection risk of group g in year y, f_g the fraction of individuals who belong to group g, $X(k)$ the survival rate into sexual activity at age k, and $S(t)$ the survival rate t years after infection. In practice, the data most resemble the unrealistic case where peak risk begins immediately, regardless of the age of the epidemic. Then, the number of deaths at age a would be

$$\begin{aligned} D(a,2002) &= POP * \sum_{g=1}^{2} \sum_{\tau=0}^{a-k} \sum_{k=0}^{a-1} \left(X(k) - X(k+1)\right) f_g * \\ &\quad \prod_{y=2002-(a-k)}^{2002-\tau} \left(1-\rho_g(y)\right) \rho_g(2002-\tau) \left(S(\tau) - S(\tau+1)\right) \end{aligned} \tag{1}$$

where POP is the population of reproductive-aged women in South Africa who are at some risk of infection. POP is a free parameter, which I choose so that the peak of the two distributions are equal; comparing the size of the at-risk group to the population generates

the size of the third group. To insure that the high risk and low risk group are different, I allow ρ_1 to reach a peak incidence of between 0 and .2 (.2 is the highest considered as this is the annual risk from unprotected sexual activity with an infected spouse among highly sexually active young couples (Gray et al 2001)), while I restrict ρ_2 to be between 0 and .02 to create a low-risk group. 2% per year is a high level of "low" risk; it creates more than 50% lifetime chance of infection. I further assume that $\rho_g(y)$ increases linearly from the first year of the epidemic until it reaches peak risk, at which point it remains at that level until 2002. Since I do not know which year peak incidence would have been reached or what peak risk should look like for each group, I allow the model to choose whichever year of peak risk fits the data best, and whichever peak risk rate similarly fits the data best, subject to the above restrictions. To give this model the best possible chance of fitting these few data points, I also allow it to choose f_1 anywhere between 0 and 1. To perform this calculation, I specify a relatively fine grid across the variables (ρ_1, ρ_2, f_1, and the year of peak incidence), calculate the implied age-death profile, and search across them for the solution which minimizes the sum of squared deviations. I further examine the epidemic at a variety of ages, given all of these assumptions. 15 years seems about right for South Africa in 2002– the relatively flat age death curve in 1996 suggests that few could have been infected before 1987.

As it turns out, in all specifications, year 1 is chosen as the best fit, meaning that the age-death curve most resembles the unrealistic situation where peak incidence was achieved immediately. Figure 5 shows that a 35 or 45 year old epidemic can fit the age-death profile quite well – the intuition that, if everyone at risk is at high risk, they attrit from the population at young ages holds. However, when the infection is young, it cannot. In fact, around year 15, the highest risk is amongst the oldest groups, in contrast to what we observe in the data. Preferences and other age-independent heterogeneous risk models are able to generate a similar curve to the one we observe when nearly all who will be infected are infected and taken out of the risk pool when they're still young – when the epidemic is young and risky people were safe in their youths, or when many people have low or middle

risk levels, it provides a poor approximation.

4.2 Marriage as a risk factor

Given that age-independent models of heterogeneous risk are unable to generate the decline in risk which we observe in South Africa, we can be certain that age-increasing models are poor approximations to the behaviors by which most people contract HIV. As argued above, this suggests that the presence of other sexually transmitted diseases do not serve as a stand-alone story for the age-death profile in South Africa. Another story which suggests age-increasing risks is a popular story about the implications of marriage for HIV, which describes how wives are subjected to the continued risk that their husbands assume through extra-marital affairs. For example, after Nicholas Kristof interviews several African women who are certain that their husbands have other girlfriends, he writes "The stark reality is that what kills young women [in Africa] is not promiscuity, but marriage. Indeed, just about the deadliest thing a woman in Southern Africa can do is get married" (2005). However, since older (and more often married) women are even safer in South Africa than they would be if marriage had a neutral effect on risk, we can rule out the cheating husbands story as being the mechanism by which most people become infected.

This story is so pervasive that it is worth reviewing the evidence which has led to it. There is clearly some risk associated with a long term partnership, for both men and women – de Walque (2007) documents that in a large fraction of infected couples, only one member (as often the woman as the man), is infected. However, a lack of long-run panel data, combined with mortality-related attrition and low transmission rates of HIV has left us without a good sense of how that risk compares to the risks of pre-marital behaviors. Excepting de Walque's paper, much of the evidence on extra-marital risk has come through anecdotal and ethnographic channels, which are hard-pressed to answer questions of relative magnitudes. Even if a large fraction of adults have extra-marital partners, risks may be minimized if the existence of a spouse prevents high-frequency sex with these partners. Empirical evidence

on marital risk has been mixed, with Clark (2004) finding that married teenage girls have higher prevalences than their unmarried peers, concluding that marriage is risky. In contrast, Bongaarts (2006) finds that marriage itself is less risky than being sexually active and single, which is supported by Glynn et al's (2001) finding that married women who avoided premarital sex have lower prevalence than their married peers who did not. These findings can be brought into accordance with each other and the Age-Death profile through an endogeneity story: spousal search may itself be responsible for much of the spread of HIV. In what follows, I specify a model of spousal search, verify that it can create an HIV epidemic with an age-death profile like that observed in South Africa, and contrast data on HIV, singlehood, and marriage across Southern and Eastern Africa to model predictions.

5 Model

The model is an adaptation of Jovanovic's (1979) model of job turnover. Individuals live for T periods. Each period, individual i receives utility $\theta_{ij} + q_j$ from a match with partner j, where $\theta_{ij} \sim F(\theta)$. q_j is the observable component of quality, and θ_{ij} is match-specific and unobservable. She faces a choice at the end of the period: to stay with partner j or to draw partner j', with whom $\theta_{ij'}$ is unknown but $q_{j'}$ is known[6]. Match quality evolves stochastically according to $H(\theta)$. Individuals match assortitatively on the q_j; in practice, this observed element drops out of any individual's optimization problem as it is constant among any mutually acceptable suitor. Therefore, each individual solves the dynamic programming problem with Bellman equations

$$V_t(\theta_{ij}) = \max_{\{stay,leave\}} \left\{ \theta_{ij} + q_i + \beta E\left[V_{t+1}\left(\theta'_{ij}\right)|\theta_{ij}\right], E\left[\theta_{ij'} + q_i + \beta V_{t+1}\left(\theta'_{ij'}\right)\right] \right\} \quad (2)$$

[6]Individuals in this model draw only one partner at a time. While age-independent concurrencies perform poorly against the age-risk profile, search with multiple concurrent draws would have similar predictions to this model, with even more severe epidemic prevalence implications. Policy predictions, however, would remain similar.

where β is the discount rate and θ'_{ij} is the random variable representing next year's match quality with partner j. As is well known, the solution to this problem is a sequence of reservation qualities, $\bar{\theta}_t$, where (in this discretized model) $\partial\bar{\theta}_t/\partial t \leq 0$, and where individuals stay in any relationship where $\theta_{ij} > \bar{\theta}_t$. Individuals in this model usually have clusters of a few very short relationships in between a few longer-term boyfriends and girlfriends. As they age, reservation qualities lower so that the probability of a good relationship turning sour enough to dissolve after a certain age becomes low – this fact creates "marriage" behavior. In my preferred specification, match evolution is slow, so that the first period provides a fairly accurate measure of the quality of a relationship; this evolution rate is chosen to mimic the observed 9 year time-span from sexual initiation until marriage in South Africa and keep the number of partners down to reasonable levels. This evolution can be interpreted as truly evolving utility from a partner or as a learning process. I define marriage as occurring at the date when a person begins to match with his (ex post) last partner. For simulation purposes, a new cohort enters every four months. Men and women match randomly with someone else who is actively searching for a new partner. After 20 years, individuals quit searching, with a payoff of 10 years' worth of utility at the current match quality[7]. Men and women are identical in this model; this assumption is supported by data presented in de Walque (2007), who observes that both women and men are the only infected member of a large fraction of couples with only one infection.

5.1 Integrating HIV into the matching model

A simple model would focus on the average rate of infection per partnership year. This has been well documented through longitudinal studies in Uganda and Zambia. By following

[7]In a more sophisticated model, individuals may consider the state and future path of the HIV epidemic in making their spousal choices. However, if individuals have rational beliefs about HIV, the distortion in behavior is minimal. As the risk from drawing a new partner is extremely close to the risk from staying with the current partner when you are still early in a relationship, individuals are only very marginally willing to lower their reservation in response. Given that the rational difference is extremely small, and that I have no way to assess what South Africans actually believe about the transmission rate or the future path of the epidemic, I abstract from this analysis.

sero-dischordant couples (where one member is HIV-positive and the other is HIV-negative) through time, medical researchers are able to observe infection rates per year and sometimes per contact. The studies in the Rakai district, Uganda, are particularly compelling as participants report minimal condom use despite counselling. Here, average infection rates of about 12% per partnership year (PPY) have been observed (Gray et al 2001, Quinn et al 2000), which corresponds to an approximately 1 in 1000 infection rate per contact. Transmission in Lusaka, Zambia, where higher condom usage was reported, was observed to be somewhat less (8% PPY) (Fideli et al 2001).

However, the average infection rate may not be a sufficient statistic to understand the dynamics of HIV transmission. PPY rates do not describe within-individual or between-individual heterogeneity in infectiousness, and there is a great deal of evidence that both are quite important for HIV. Gray et al (2001), Quinn et al (2000), and Fideli et al (2001) all find viral load in the blood plasma has a very large effect in predicting HIV transmission in the African setting, confirming studies from the US which have found the same result (e.g. Lee et al (1996), Ragni et al (1998), Pedraza et al (1999)). That is, some individuals have much higher viral load than others, and it is precisely these individuals who are most likely to infect their partners. In terms of viral load, we can divide a person's HIV infection into three broad periods. First, acute infection lasts for the first two to three months. The body has not yet developed an immune response to HIV, and viral load is high. Next, in latent infection (the next eight years or so), the body's immune response keeps viral load extremely low. Finally, the body's immune system starts to lose the battle, viral load climbs again, AIDS breaks out shortly and within a year or so the individual dies without medical intervention (e.g. Katzenstein (2003) for a review). People in the longitudinal studies described above are largely in these second and third phases, and individuals in the third phase are unwell enough that high coital frequencies seem unlikely, which is why low transmission rates are consistently observed. Pilcher et al (2004) note that infectiousness during acute infection may be 10 times as high as in latent infection, a number confirmed

empirically by Wawer et al (2005). In the epidemiological model most similar to this one, Koopman et al (1997) assume individuals transition randomly between high and low-turnover states and match non-randomly with others in the same state. Considering the US, and using the larger transmission probabilities associated with anal sex, they find that shutting down acute infection may end the epidemic entirely.

Therefore, in simulating the model, it is important to take into account the dynamics of HIV infection. To do so, I specify that if an individual of gender g matches with a partner is in acute (latent) infection, they face probability ρ_1^g (ρ_2^g) of infection. I ignore the last stage of infection, as individuals in this stage are ill and may not be maintaining high coital frequencies despite their greater viral load. Finally, after nine years of infection, I "turn off" infectiousness for individuals; nine years corresponds to the mean survival length for HIV in Africa. I do not allow individuals whose partners have died from HIV to find a new match, unless they would have resumed searching from the relationship's evolution in any event. In each cohort, a small percentage (δ) enter already infected with HIV, with $\delta/6$ entering in acute infection[8]. This consistent injection of HIV allows me to obtain higher prevalences and can be interpreted as infection from all other sources. Individuals who are initially infected do not behave differently from other individuals.

Each period corresponds to a month, and mature transmission rates are similar to those found by Gray et al (2001) for young couples (slightly above average rates are due to the greater coital frequency of young couples) whereas acute transmission rates follow Pilcher et al (2004) and Wawer et al (2005). The distribution over θ is chosen for simplicity; changing to other simple distributions alters lifetime numbers of partners and HIV prevalences surprisingly little, as reservation qualities adjust downwards when good matches become more scarce. Men and women are identical in this model, and as such have very similar simulated data; I present only the results for simulated women.

[8] 1/6 was chosen as I find from simulations that about 1/6 of single, HIV-positive individual-months are spent in acute infection for those who don't enter already infected. Varying this fraction has only a small impact on results.

5.2 Simulating the Model

To simulate this model, I first solve the dynamic programming problem numerically according to the distributional assumptions given in table 4. Then, I simulate N men and N women in the first cohort, each of whom randomly draws a partner from the simulated individuals of the opposite sex and receives a draw of partner quality. At the end of the period, they decide whether to stay in the relationship or draw a new one based on the solution of the dynamic programming problem (if either individual prefers to end the relationship, it ends). If they stay, their match quality evolves and they face the same decision in the subsequent period. If they leave, they receive a new partner from the pool of other individuals who choose to search and a new draw on θ. Every ten months, a new cohort of N individuals of each sex enters the pool of searching singles and similarly draws a partner and a match quality. I simulate the model for 40 years; I do not explicitly model homophily so that entering cohorts are treated the same as single individuals in existing cohorts (though of course their age is different, resulting in a different solution to the search problem). Table 5 features summary statistics from 2500 observations of the simulated data. With the entry of the 10th cohort, the HIV epidemic enters, and all living individuals face instantaneous probability δ of contracting HIV[9]. Following, each entering cohort enters with starting prevalence rate δ as described above, and HIV transmission occurs as HIV-negative partners of HIV-postive individuals face random draws against the transmission rate. There are 2 sources of randomization in this this model: the match process and the HIV epidemic. Once the match has been solved, the spread of HIV happens in a way which is relatively quick computationally. All results

[9]Since cohorts must build up over time, HIV entry is delayed to minimize the effects of time dynamics in homophily (simulations suggest that this time-delay has a slight and conservative effect on HIV prevalences). In fact, the role of homophily in the spread of HIV is somewhat unclear – if people match non-randomly with others in a similar age-category, this may lead to more explosiveness in the HIV epidemic as people in high turnover phases with greater likelihood of acute infection match systematically with others in the same phase. This observation makes the role of homophily in the spread of HIV an important avenue for future research, particularly as it interacts with "relative risks" messages (Dupas 2007) which persuades girls that older men have higher prevalence (and hence that younger boyfriends are safer). Fortunately, the lower pregnancy rates found by Dupas (2007) for girlfriends of young and presumably high-turnover boys suggest that lower fertility demand may overwhelm this concern.

presented are for 50 realizations of the match, each averaged over 10 realizations of the HIV epidemic.

6 Results

Table 6 reports the fraction of fully exposed women who will become infected for each δ. Every percentage point of infected individuals which I input into the model results in about 6 additional infections. In other words, each person who enters the spousal search period infected will cause about two and a half other single people to become infected, and all of them will go on to infect the person they eventually marry. Thus, if 1% of the population enters the spousal search period infected, we end up with a prevalence rate similar to that in Kenya or Tanzania. With just 3% receiving infections from other sources, we are at South Africa's very high epidemic prevalence rates.

The infection moves quickly; within the lifetimes of the first six cohorts, the prevalence rate reaches around 80% of the level that it ever will. By the end of the epidemic, each infection is responsible for about 7-8 additional infections, with some concavity with respect to the inputted level. The conclusion of this section is strong: with just a tiny fraction of the population entering the spousal search phase of their lives already infected, marital shopping can create an epidemic, and create one quickly with relatively few partners. The reason for this is basic behavior generated by matching models (and hence fairly insensitive to parametric assumptions). Individuals have clusters of several very short relationships which are easily rejected in between a few much longer ones. Therefore, when they enter a new relationship, they are likely to have just left another very short one. In turn, this means they are much more likely to have just been infected, remain in the acute phase of infection, and hence more likely to infect their new partner.

One advantage of a simulation model like this one is that we can attest to the importance of acute infection by simply shutting it down. That is, I can just set $\rho_1 = 0$, and explore

what happens to population prevalence. Columns 3 and 4 of table 6 reports the prevalence results of this situation. The results are strong: three months of protection cuts the epidemic prevalence rate in half. Now, δ is multiplied by 3 rather than 7, so that each initially infected person only causes one other single infection before the two of them go on to infect their spouses. This thought experiment is interesting for two reasons. First, it supports data from Switzerland (Yerly et al 2001) which has shown that over 1/3 of infections are attributable to acute infection (Wawer et al 2005 report similar findings in Uganda). Secondly and more importantly, it is implementable. That is, this thought experiment corresponds to a policy of wearing condoms for the first three months of each relationship. The idea of a preference for condoms in short relationships sounds familliar and intuitive; it is very similar to public health messages which were spread following the first reported cases of HIV in the US. However, it represents a sharp divergence from the current message being spread by many NGOs, who promote condom use throughout marriage (e.g. Ali et al 2004). In the policy discussion section, I will evaluate the consequences of this policy in this model more explicitly.

Does the matching model generate an age-death profile similar to the one observed in South Africa? Since I am interested in confirming the plausibility of this story rather than rejecting it, my approach is very different from the one used in the heterogeneous risk case. That is, ultimately I am fitting a rather small number of data points here. It would be unimpressive if I was able to search through many parametric specifications in order to find a curve which approximated the age-death profiles for men and women in South Africa (despite the failure of the heterogeneous risk model at this task). Therefore, I rely as much as possible on theoretical predictions common to matching models of spousal search, and use data to derive the end result of the true model. Specifically, I observe the following: in a matching model, infection risk is constant for the uninfected single. When married, in contrast, it declines exponentially – since no new infections are brought into a marriage, your risk at any point in time is the product of the transmission rate, the risk that your spouse

entered the marriage already infected, the chance that you entered the marriage uninfected, and the probability that you haven't caught it from your spouse yet. Similarly, the time path of risk suggests the relative fraction of infections occuring to the single and married. The algebraic identities suggested by these predictions are straightforward if tedious, and are contained in the appendix.

The South African DHS provides empirical distributions of the age at first marriage and the age of first sexual activity; I estimate Kaplan-Meier hazard functions into marriage and into sexual activity to predict these, with the assumption that being unmarried but having had sex is reflective of being single and searching[10]. As the DHS only provides data on women, I predict the age of first marriage for men by taking means by age of the fraction married from the September 2001 South African Labour Force Survey, and I assume either that the hazard function into sexual activity is the same for men as it is for women, or alternatively that it is shifted to five years older as South African men marry five years later than women on average. As described above, I know the risk profiles for the single and for the married if I know the time path of the epidemic and the relative prevalence rates for newlyweds and singles. Therefore, to construct the age-death profile, I need only the time paths of the prevalence for single people and for newlyweds. Each of these are taken from simulations; reassuringly, the figures below are robust to reasonable changes in both of them as it is the empirical distributions of singlehood and marriage as well as the theoretical risk predictions which produce the close fit.

Figures 6 and 7 represent this comparison. Here, I make the further assumption that the number of deaths by age in 1996 represents the number of non-AIDS deaths in 2002, and I set the peak of simulated AIDS deaths equal to the empirical peak. That is, the

[10] These create an integer problem as you are only sexually active for some fraction of the year in which you first have sex. Because the age at which you first have sex is the absolute earliest you could start risky search, I assume that young adults begin searching (and hence reach the full single risk level) in their first full year of sexual activity. However, the potential for heightened pre and post-nuptial coital frequencies suggest that married individuals may start on the married risk rates in the year in which they get married; the figures presented reflect this assumption. Resolving the integer problem in different ways results in very similar fits.

height of the curves is rigged; however, that is the only point which is expressly fit and the shape of my predicted age-death curve is derived from the theory and the data. As the reader can observe, the fit is very close. I very slightly underpredict older women's deaths and overpredict younger ones. Recalling that this model assumes perfect monogamy, a small amount of extramarital behavior may create this difference; however, the quantity of infections that this non-monogamy produces is tiny in comparison to the number of monogamous infections. For men, the data lie in between the age death curves created by assuming that men have the same sexual initiation as women and that assuming that their sexual initiation is five years older. Still, the fit is overall very close, and distinguishable from the fit for women. Men marry five years later than women and over a broader range of ages, and that is precisely how they die.

In evaluating the closeness of fit, I propose two metrics. The first, given in figure 5, is the most successful that a heterogeneous risks model can achieve. According to this metric, the matching model of serially monogamous spousal search emerges clearly preferable. It is also preferred to the current state-of-the-art epidemiological modelling for South Africa, given in ASSA (2005). The ASSA model embeds a system of a large number of (non-behavioral) equations into an effort to predict several characteristics of the HIV epidemic, including the age-death profile.[11]. Figures 8 and 9 reports the fit generated by the ASSA model. The single equation of the matching model seems to outperform this model rather substantially for both men and women, with average sums of squared deviations of about half the magnitude.

There is no reason to expect that the marital shopping model would only fit the epidemic in 2002. Fortunately, we can estimate the model for every year, 1997-2002, presented in figures 10 and 11. For men, recalling that the true data appeared to be halfway between the two assumptions on sexual behavior, I take the average of the predicted results from these two assumptions. For both men and women, at all years of the epidemic, we find a close fit

[11]Notably, marriage is not considered in the ASSA model.

(once again, only the peak of each year is fixed). I consistently underestimate prevalences at the oldest ages by a small amount – these may well be the individuals who enter in δ. Moreover, some of the time dynamics in the age-profile of infection are preserved between the model and the data, not least being the age-peak from early years which was absent in the heterogeneous risk model. Regarding the female age-prevalence profile, we see that at the onset of the epidemic, peak deaths occured in the 25-30 age group, with a similar rate among 30-35 year olds, whereas by 2002 peak deaths were the sole province of 30-35 year olds. This is predicted by the marital shopping model.

7 Empirical Support

Thus far in this paper, I have evaluated existing behavioral epidemiological models of HIV and found that they do not fit the empirical age-trend of infection in South Africa. In contrast, I have proposed a matching model of HIV which is capable of generating these high prevalences and suggests an age-pattern of infection very similar to the observed one. This section argues that the matching model is also consistent with micro-level data on HIV across Southern and Eastern Africa. One of the challenges in doing so is the (much) greater prevalence both of concurrent and of consecutive marriages in the other countries of sub-Saharan Africa. For example, only 14% of South African women over 30 are divorced, in polygamous relationships, or in a marriage other than their first (with younger women being less likely). In contrast, 28% of Zimbabwan and 46% of Malawian women over 30 belong to one of these groups (from the DHS data). The spousal search model is easily adjustable to accomodate these, the first through offering multiple search draws and the second through relabeling "marriage." However, it would require strong modelling assumptions to generate analogous theoretical predictions on single and married risk to those used earlier (for example, are polygynous men constantly searching? Or do they simply receive an additional spouse at some point?). Thus, the empirical focus in this section is on implications of spousal

search which are less sensitive to these sorts of assumptions: in particular, that marriage is safer than singlehood for non-polygynous individuals whose spouses enter the marriage uninfected[12], and that a longer time period in spousal search is associated with elevated risk.

If spousal search carries much of the risk, then couples who have been married for several years and who are still alive should have substantially reduced risk. That is, if at least one member of the couple was infected before marriage, it is unlikely that both will remain alive ten years into the marriage. Some evidence to this end is presented in Beegle and de Walque (2008), where, in all the high prevalence countries available from the DHS, mean prevalence rates decline when comparing couples married over ten years to the sample average. To get a good estimate on the size of this effect, a first pass approach is to look at the probability that at least one member of a couple is HIV-affected and view how that varies with marriage tenure. Table 7 estimates these correlations for the high prevalence countries in Southern and Eastern Africa with publicly available DHS data[13]. Analysis is restricted to monogamous couples (except in Lesotho where the polygamy question was not asked), as polygamous individuals seem likely to have HIV-risk vary with marriage tenure in a different way than monogamous couples. The estimation procedure is a probit with five year age-group and spousal age-group fixed effects used to correct for the correlation between years of marriage and age. The dependent variable is equal to one if the couple is unaffected by HIV, that is, if both members test negative, and throughout the omitted group are couples who have been married 0-5 years. In every country, couples who have been married for at least ten years are substantially less likely to be infected than their recently married peers, with precisely estimated coefficients for every country save Lesotho. Ten years of marriage reduces your likelihood of infection by 50%-100% of the mean prevalence level, relative to peers who have

[12]Even this prediction is somewhat ambiguous if non-polygynous individuals are simply individuals whose spouses are still searching for their second wife; to the extent to which this is true it should bias against the results found below.

[13]Summary statistics for these countries are in the appendix. All of these countries had HIV-test refusal rates between 14 and 20%; all analysis is subject to that caveat.

been married 0-5 years. Kenya makes for an extreme example. Among the 486 couples married less than 10 years, 65 of them have at least one member infected, whereas only 24 of the 411 couples who have been married at least ten years have at least one member infected. Of these, in only 14 (5 women and 9 men) is one member infected without their partner being infected; since a few of the jointly infected couples will be long-term survivors from pre-existing infections, we can guess that somewhere from 14-24 infections were brought into these couples from outside the marriage. Though endogeneity may hinder causal interepretation of this analysis, this analysis is consistent with the predictions of spousal search as a first-order risk mechanism.

A second prediction of the marital shopping model is that length spent single should be correlated with HIV infection. That is, a longer time period spent single indicates less luck in finding a match, and should be correlated with the amount of risk borne. Bongaarts (2006) suggests empirically that both age at marriage and the gap between the age at first sex and age at marriage are important predictors of HIV infection. Similar to Bongaarts's analysis, table 8 predicts HIV prevalence at the sample cluster and individual level using the DHS surveys. Here, I estimate

$$\bar{Y}_i = \beta_1 \left(\overline{Age1stSex_i}\right) + \beta_2 \left(\overline{AgeMarried_i}\right) + \beta_2 \left(\overline{Polygamy_i}\right) + \beta_3 \left(\overline{AgeDiff_i}\right) + \varepsilon_i$$

Where $\bar{Y}_i$ is the HIV prevalence in sampling cluster i, $AgeDiff_i$ is the difference in age between spouses, and variables represent sampling cluster-level means. If the length of singlehood is important, then we should observe that the effects of the age of first sex and age of first marriage are opposite in sign and similar in value. Odd columns reveal that later marriage and earlier sexual onset have similar effects on HIV prevalence at the sample cluster level and that some characteristic about time spent single is indeed correlated with HIV prevalence at the population level across countries. These results are made more striking by the fact that two other likely candidates for HIV risk, polygamy and spousal age

differences, are not consistently statistically significant[14].

A natural extension is to examine whether the correlation is present at the individual as well as at the sampling cluster level. Even columns of table 8 restrict attention to non-polygynous women in their first union, and condition on HIV prevalence of other women in their sampling cluster. These columns reveal that this same characteristic is also correlated at the individual level across countries. However as both the decision to commence sexual activity and the decision to get married are endogenous, we may remain concerned that this the length of singlehood is picking up other correlated omitted variables like preferences. I address this issue using the insight that infections from pre-marital behavior should disappear with marriage tenure. That is, if singlehood is risky, then the effect of the years single should be strongest over the recently married, as those who have been married longer are unlikely to have survived pre-marital infections. In contrast, if a long period of singlehood is simply a signal of preferences, then those with long singlehoods should still be at risk years after marriage. Table 9 presents marginal effects from a probit of HIV prevalence on years of singlehood and dummies for each marital tenure category, where the effect of singlehood on HIV is allowed to change with marital tenure. Across countries, years of singlehood are most strongly related to HIV status in the first ten years of marriage; in all countries, the point estimates go down substantially after these years (and for Lesotho and Zimbabwe, we can statistically distinguish the effects of years 0-5 from years 15+ at at least the 5% level). Ten years is a logical turning point, as it is the median life expectancy after infection. These results suggest that being single longer is risky in large part due to behaviors which take place when single rather than correlated behaviors which last a lifetime, consistent with the marital shopping model.

Another prediction of the marital shopping model is that single people should have a large variance in their annual number of sexual partners. In particular, a year of singlehood where

[14]One concern may be that sex and marriage behaviors have adjusted to accomodate the differential HIV prevalences in sampling clusters. I repeat the analysis using sex and marriage variables for women over 40, for whom these decisions were made before there were many documented HIV cases. Once again, years single seems to be the primary determinant of HIV prevalence (results are noisier but still mostly significant).

you have a sequence of bad draws in partner quality will result in a large number of partners, whereas you may have few partners the next year with one good draw. Models based on heterogeneous types do not deliver this result. In, for example, Kremer's (1996) specification, if we observe an individuals with many partners it is a signal that that individual is a high-risk type, and we should expect him to have many partners again in the following year. In fact, if we know the number of partners in one year, and the number of years of sexual activity, we can perfectly predict the lifetime number of partners in a heterogeneous types models by simply multiplying the annual number of partners with the years of sexual activity, subject only to integer problems[15]. Both Cape Town and U.S. data prefer the matching model to a preference model in terms of a Vuong (1989) model selection test (results available from the author).

Given that both the African and American data prefer the matching model, it is interesting to ask just how similar they are. Columns 1 and 3 of table 3 report tabulations of numbers of partners by gender in the last year using U.S. adolescents from the NSFG, while columns 2 and 4 report similar calculations using the CAPS teenagers who live in Cape Town (all data is described in the data appendix). Across cultures, the annual numbers of reported partners appear identical. To the extent that adolescent males are only partnering with adolescent females, we know that one of these two groups is misreporting. Strikingly, even the bias in male versus female reporting appears to be consistent between the two continents. Given similarities between the US and Africa an additional test of the model would be to consider whether inputting transmission dynamics of a different sexually transmitted disease in a different context also provides a good age-fit. Unfortunately, transmission probabilities of most sexually transmitted diseases are little understood and cases are often undocumented. Gonorrhea in the US, however, provides a good case study.

[15] We may be concerned that the integer problems are especially severe for people reporting one partner per year who may in fact be having far fewer. I find that assuming that individuals with one annual partner have one lifetime partner fits reported sexual behavior data far better than assuming one annual partner, and so adopt this assumption in my specification. The Cape Town data Source is CAPS, while the US source is the NSFG, both described in the appendix.

Unlike HIV, gonorrhea is extremely infectious, with transmission probabilities very high for a single contact and approaching 1 for a month-long relationship. Gonnorhea is also a transient infection, with most people spontaneously recovering without treatment in a few months, or experiencing quicker recovery with an anti-biotic. The marital search model, then, would predict a constant incidence for single, sexually active adults and zero incidence for married adults. Figures 12 and 13 illustrate the predicted versus observed gonnorhea prevalences by age in the US. For both women and men I overpredict prevalences at older ages. Nonetheless, the predictions do exhibit a similar pattern to the data despite the very different biological and geographical context.

8 Policy Discussion

This paper considers two policies to reduce the spread of HIV. The first, suggested by evidence on the importance of acute infection is for individuals to use protection in their short relationships. This eliminates the role of acute infection in the epidemic. However, we may expect that, under the framework of marital shopping, using condoms in short relationships could be even more effective than shutting down acute infection. The vast majority of relationships in this model are short in tenure, as most partnerships are quickly rejected as unacceptable, which means that using condoms at the beginning of relationships may protect you from any transmission from most of your partners. At the same time, it may be impossible to get everyone to wear condoms in their short relationships; as such, I consider what happens when only some fraction λ of the population protects itself for the first three months of each relationship. Table 10 observes the mean lifetime prevalences created across cohorts for a variety of λ,in each case assuming $\delta = .03$, which relates to the South African epidemic[16]. For each value of λ, the overall prevalence is reported and then decomposed into the prevalence for those who follow this strategy and the prevalence for those who do not.

[16] Other $\delta's$ have their multiplier reduced similarly under this policy and are omitted for space considerations

The strategy of wearing condoms in short relationships is very effective for those who follow it – in each case, the inputted $\delta's$ are multiplied by 3-4.5 for the people following this strategy and 4-6.33 for those who do not. We see strong evidence of externalities here; as some individuals protect themselves, total epidemic progression is strongly limited, and everyone faces lower risks as a result. Moreover, if everyone uses protection in short relationships, we see a tremendous fall in epidemic prevalence – now, each new infection is responsible for about 2 additional infections. Given that, in a model like this, spouses will almost deterministically become infected, that means that each inputted infection causes about half of an additional couple to become infected (rather than 2.5-3 couples in the baseline case).

Much as the motivation for the behavioral model discussed in this essay builds upon Kremer's (1996) work, the policy predictions and simulated evidence developed here are similar in spirit to Oster (2005), who argues that the differential extent of the HIV epidemic in Africa can be attributed to differences in transmission risk due to greater prevalence of secondary sexual infections. Oster computes stark differences in the transmission risk between Africa and the United States, finding that African transmission rates are three times as high as those in the United States. The policy prediction from such a finding is clear: African transmission rates should be reduced to US levels. A similar policy is prescribed by Potts et al (2008), who emphasize that male circumcision may be the most feasible way to bring African transmission rates in line with US levels.

Using the marital shopping model, we can observe the effect on prevalence of changing ρ_1^g and ρ_2^g to reflect differences in transmission rates, while remaining agnostic about the best mechanism to achieve such a result. As such, it is appropriate to model how much of a transmission rate change is feasible biologically. As discussed above, transmission rates in Africa range from 8-12% per partner year. US and European studies have found PPY transmission rates in the 5-10% range (see column 8 of table 11)[17]. This constrasts sharply with the large differences in Oster (2005). In fact, the epidemiological data used to derive

[17]This similarity in the face of lower condom use, more circumcision, and higher STI prevalences led Gray (2001) to reject the hypothesis that HIV-1 subtypes in Africa are more infectious.

transmission rates put forth in this paper are many of the same sources as those used in Oster, but the finding of vastly different transmission rates is not preserved. This is because of differences in how transmission rates are measured. Oster follows papers such as Downs and Di Vicenzi (1996), who propose using a per-partnership transmission rate rather than a per-contact rate, as heterogeneity and dynamics in infectiousness leads to an average infection profile which is not binomial in sexual contacts. This approach represents one approximation; it is imperfect as individuals do become infected several years into the described studies, suggesting that additional coital acts do indeed increase your transmission risk[18], a limitation which seems especially severe in describing transmission rates in mostly short relationships. The variance in transmission rates across measures emphasizes the quandary in choosing the right one, summarized in Table 11. It appears that no matter how transmission rates are estimated, a wide variety of point estimates are received, with per partnership rates exhibiting a particularly high variance. This is due to factors such as differences in lengths of the studies, differences in coital frequencies, differences in the fraction of study participants who report perfect condom usage, small sample sizes (particularly for female to male transmission), and of course different biological transmission rates due to STIs, circumcision, and different viral strains. Notably, none of the African studies find statistical differences between Male-Female and Female-Male transmission despite widespread belief otherwise. The more recent studies, all African (as ARV availability in the developed world limits the potential of such a study) choose to report per-partnership-year or per-coital act transmission rates, which appear to report more consistent transmission rates. Both of these measures suggest that transmission in the US and Europe is on the order of 50%-100% as efficient as in Africa, so I take the midpoint and assume that a 25% reduction is achievable through biological means. Given that participants in the Western studies report perfect condom usage more frequently, we may well expect them to have more

[18]In fact, a depressing reality in many of these studies is that many people become infected only after their partners viral loads begin climbing with the transition from HIV into full-blown AIDS, at which time survival rates are vastly diminished.

consistent imperfect use; this suggests that the 25% estimate may be conservative.

A one-fourth reduction in transmission rates may still have vast epidemiological consequences. Rows 1 and 2 of Table 12 report the results of repeating the marital shopping model, with transmission risks in both acute and latent infection at 87.5% and 75% of the previous levels. The 87.5% exercise represents moving halfway from the African transmission rates to the US ones; this appears to reduce the multiplier by about 1. Moving all the way to US transmission rates changes the multiplier property from between 6 and 7 to 4 – now, rather than each inputted causing two and a half to three additional couples to end up infected, each causes one and a half to two. Clearly, this is an effective strategy, about as effective as compelling half of individuals to use protection in their short relationships; which of these approaches is more feasible is an empirical question. The remainder of table 12 reveals what happens when transmission rates are reduced in parallel with individuals using condoms in short relationships. The results here are strong. In this model, there is basically no way to reduce the multiplier below two – individuals who enter spousal search already infected will with very high probability infect their eventual spouse. When we reduce the transmission rate to American levels and compel individuals to use protection in their short relationships, we find that very few individuals other than spouses ever become infected. If Americans have these transmission rates and follow this strategy of protection in short relationships, it may explain why spousal search does not create an HIV epidemic in America despite otherwise similar sexual behavior. Strategies which seek to reduce transmission rates through treatment of sexually transmitted diseases and male circumcision can effectively combat the epidemic, as can short intervals of protected sex, and the pursuit of both strategies in tandem may be tremendously effective.

9 Conclusions

Serial monogamy with high turnover is sufficient to create and maintain extremely high prevalence levels, and can blow up infection from other sources to much higher levels. Simple dating can create this behavior, if we believe that there are idiosyncratic, unpredictable components to the quality of a relationship and that individuals prefer spending more time with better matches. The age-profile of deaths is extremely restrictive as to which explanations for the spread of HIV it will allow; the matching story stands above many existing models in passing this test. Moreover, the matching story allows acute infection to be important without individuals having extremely high numbers of lifetime partners, supporting empirical evidence on the importance of acute infection.

If individuals use protection in only their short relationships, prevalence rates fall dramatically, with 50% usage being as effective as a reduction in transmission rates all the way to US levels. This highlights an important choice for public policy: should policy makers emphasize using condoms in new relationships at the cost of using them in older ones? Some individuals may find 3 months of condom use much more palatable than a lifetime of protection, and a tremendous amount of risk would be averted[19]. Indeed, this message contrasts strongly with the message adopted by many public health groups, who encourage condom use throughout marriage. As many individuals doubtless hope to have unprotected sex at some point in their lives, it would be truly dangerous if they felt they had to make an all or nothing choice because they did not understand the relative risks of pre-marital and marital sex. Moreover, when short-term condom use is combined with reduced transmission rates, marital shopping ceases to be an effective mechanism to transmit HIV, suggesting two clear goals for policy.

Like any advertising campaign, public health campaigns are targeted at specific groups. My analysis suggests that the most important group to target is young and single men and

[19] Adolescent girls in rural Kenya have been shown to respond to information on relative risks of different partners by changing their sexual behavior(Dupas 2007), so this message of targetted risk seems likely to be one that individuals can understand.

women, and the correct message would indicate that single, monogamous relationships carry a great deal of risk, particularly in the first months. This is in contrast with campaigns whose primary message is to encourage monogamy, or those who blame extreme-risk behaviors or individuals for the spread of HIV. It would be dangerous if individuals believed that monogamy implied safety; in particular, every relationship is monogamous on the first date, when the probability of acute infection is highest. Moreover, widely-held beliefs demonizing married men for HIV seem inconsistent with the observed deaths profile; these beliefs may increase HIV stigma with little empirical basis.

An implication falling out of the medical literature is that testing campaigns will be hard to sell. Antibody-based tests are by far the cheapest and those predominantly in use both in Africa and in the United States. These tests cannot pick up acute infection, and therefore misdiagnose HIV when it's at its most infectious, for the simple reason that the body has not yet developed an immune response. Since half of the risk that a person faces with a new partner is caused by acute infection, demanding that a partner get tested before intercourse will not protect the person very much – moreover, if a negative result encourages choosing against condoms, then it could actually make sex much more dangerous. If individuals learn from their friends who do demand testing from their partners and nevertheless get infected, we should not be surprised at how little testing has caught on.

Finally, this paper is fundamentally about how most people become infected with HIV, with marital shopping underlying 5/6 to 6/7 of all infections. I show in this paper that serially monogamous dating and marriage behavior acts as a multiplier on very small prevalence rates from outside sources, and very small changes in the number of these outside infections can lead to huge differences in the steady-state of an epidemic. A number of papers (e.g. Oster 2005) including this one have documented the similarity in reports of sexual behavior between Africa and the United States, and this paper has reemphasized how similar behavior with different transmission rates and short term condom usage can explain the different prevalence levels. However, a few pivotal people may well behave differently

and become infected through other methods, and both the aggregate statistics which are central to this paper as well as traditional econometric analysis are ill-suited to identifying these small samples, perhaps explaining the divergence between these results and the reports of ethnographers. Respondent-driven sampling methodologies (e.g. Heckathorn 1997) may be an asset in understanding the behavior of this pivotal minority.

10 Data Appendix

Four data sources are used in this article.

1) The Cape Area Panel Study[20] (Lam et al 2006) is a random sample of 4758 young adults aged 14-22 in 2002 who live in the Cape Town Metropolitan Area. Located in the South-Western corner of South Africa, Cape Town is the second-largest city in the country. These young adults were interviewed first in 2002. A subset of 1360 young adults were reinterviewed in 2003, with the remainder reinterviewed in 2004, and all were reinterviewed in 2005. This paper utilizes information on annual partners collected in 2002 and lifetime partners collected in 2003.

2) DHS data: this refers to the 1998 South Africa, 2003 Kenya, 2004 Lesotho, 2004 Malawi, 2003 Tanzania, and 2005/06 Zimbabwe DHS surveys. DHS survey information is available at www.measuredhs.com.. Table 13 provides summary statistics from these data.

3) Labour Force Survey: This is a twice yearly rotating panel of households in South Africa designed to examine employment status and demographic variables across South Africa, with sampling done from 1996 census blocks. More information and data requests are available at the South African Data Archive, www.nrf.ac.za/sada

4) NSFG data: this refers to the 2002 National Survey of Family Growth, collected by the National Center of Health Statistics (NCHS). The survey population is designed to be representative of the United States, and more females than males were surveyed (7643 versus

[20]The Cape Area Panel Study Waves 1-2-3 were collected between 2002 and 2005 by the University of Cape Town and the University of Michigan, with funding provided by the US National Institute for Child Health and Human Development and the Andrew W. Mellon Foundation.

4928). More information is available at www.cdc.goc/nchs/nsfg.htm.

11 Appendix: Age-Death Algebra

Let $S(t)$ be the survival rate after t years of infection and $I(a,t)$ be the overall incidence of HIV at age a in time period t. The number of deaths from HIV at age a in year t is given by

$$D(a,t) = \sum_{r=0}^{a} I(a-r, t-r)\left(S(r) - S(r+1)\right) \tag{3}$$

The marital shopping model suggests that a sharp change in risk behaviors occurs at time at marriage, so to infer death rates it is convenient to partition individuals into married versus single. That is, if $I^S(a,t)$ is the fraction of individuals who are both single and infected at age a and time t, and $I^M(a,t)$ is the same for married individuals, we have

$$D(a,t) = \sum_{r=0}^{a} \left(S(r) - S(r+1)\right)\left(I^S(a-r, t-r) + I^M(a-r, t-r)\right) \tag{4}$$

Suppose $\iota^S(t)$ is the raw incidence rate for single individuals at time t, or the fraction of single individuals who become infected in that time period. Under the matching model, single individuals all behave the same, and hence are at equal risk, with one caveat. HIV is an absorbing state, and so individuals who have already been infected cannot become infected again. Hence let $i^S(k,t)$ represent the risk of becoming infected for an individual who has been sexually active for k years. Moreover, due to HIV's absorbing nature, the raw incidence rate has to be multiplied by the fraction of individuals who can still be infected, that is, $(1-\psi^s(t))$ if $\psi^s(t)$ is the single prevalence rate at time t in order to generate the true risk that a single individual faces. Then

$$i^s(k,t) = \frac{\iota^s(t)}{1-\psi^S(t)} \prod_{r=0}^{k} \left(1 - \frac{\iota^S(r)}{1-\psi^S(r)}\right)$$

and

$$I^S(a,t) = \sum_{k=0}^{a} i^S(k,t)\,\xi^S(a,k)$$

if $\xi^S(a,k)$ is the percentage of women who are single and searching at age a and who have been searching for k years.

In the marital shopping model, married individuals face an incidence rate which declines exponentially at the annual transmission probability, ρ. In particular, let $i^M(\mu,k,t)$ be the risk of infection for an individual who married μ years earlier after k years of search in period t. then

$$i^M(\mu,k,t) = \psi^m(t-\mu)\,\rho\,(1-\rho)^{\mu-1}\prod_{r=0}^{k}\left(1-\frac{\iota^s(t-\mu-r)}{1-\psi^S(t-\mu-r)}\right) \tag{5}$$

where $\psi^m(t)$ is the prevalence rate among newlyweds, which is in general different from the single prevalence due to the declining reservation quality with age. Hence

$$I^M(a,t) = \sum_{\mu=0}^{a}\sum_{k=0}^{a-\mu} i^M(\mu,k,t)\,\xi^M(a-\mu,k)$$

where $\xi^M(a-\mu,k)$ is the percentage of women who married at age $\alpha-\mu$ after having been single for k years. I assume independence between age of sexual onset and age of marriage; in the South African DHS data these are uncorrelated. Hence if $M(a)$ represents the fraction of individuals who are married at age a, $\xi^S(a,k) = (1-M(a))\,\tilde{\xi}^S(k)$, and $\xi^M(a,k) = M(a)\,\tilde{\xi}^M(k)$, where $\tilde{\xi}^S(k)\left(\tilde{\xi}^M(k)\right)$ is the proportion of single (married) individuals who have actively searched for k years. Since this is unobservable, I assume it is proportional to the percentage of women who report having had sex for the first time at age k years earlier. Search intensity while single seems likely to be different for currently married individuals – in particular, for age-a married individuals, $\int_0^a \tilde{\xi}^M(k) = 1$, whereas there is no such implication for $\tilde{\xi}^S(k)$; so if $X(k)$ is the distribution of individuals who report sexual onset at age k, then I assume that $\xi^S(a,k) = \gamma^S(1-M(a))\,X(a-k)$ and $\xi^M(a,k) = \gamma^M(M(a))\,X(a-k)$.

Therefore

$$\frac{D(a,t)}{Popn} = \sum_{r=0}^{a}(S(r)-S(r+1))* \left(\begin{array}{c}\gamma^{S}(1-M(a))\sum_{k=0}^{a} i^{S}(k,t)(X(a-k)-X(a-k-1))+\\ \gamma^{M}M(a)\sum_{\mu=0}^{a}\sum_{k=0}^{a-\mu} i^{M}(\mu,k,t)(X(a-\mu-k)-X(a-\mu-k-1))\end{array}\right) \quad (6)$$

Using DHS data, I estimate a Kaplan-Meier survival function out of singlehood for African Women in South Africa, and a similar survival function into sexual activity. For men, for whom there is no South African DHS data, tabulations of percent never married are taken at each age from the September 2001 South African Labour Force Survey, and beginning sexual search is calibrated in two ways: as being identical to the female distribution of coital onset, and as being the female distribution plus five years (as the average married male is five years older than his spouse in South Africa). The survival function for HIV is taken from UNAIDS(2002). ρ, the incidence per year of relationship, is set to .20, similar to Gray's (2001) finding for young couples, and non-AIDS deaths are taken to be identical to those in 1996. Time-paths of single incidence rates, single prevalence rates, and newlywed prevalence rates are simulated with the model, allowing identification of everything but γ^S and γ^M. In other words, at time t, we know the shape of deaths from infections which the married incur and the shape of deaths from infections incurred by the single but not how to weight those curves in adding them. At time t, we also know the ratio of the total deaths from married infections to that of single infections. That is,

$$D^{m}(t)=\sum_{p=1}^{t-1}\sum_{j=p}^{t-1}\mu(p)\psi^{m}(p)*(1-\psi^{m}(p))\rho(1-\rho)^{j-p}*(S(t-j)-S(t-j+1))$$

Where $\mu(p)$ represents the number of individuals who are married in year p. In turn,

$$D^{S}(t)=SandS*\sum_{j=1}^{t}\iota^{S}(j)(S(t-j)-S(t-j+1))$$

where $SandS$ identifies the number of single and searching individuals. If search lasts on average nine years, as in South Africa, then this corresponds to about 8% of single and searching people being married per year. Therefore, $\mu(p) \cong .08 * SandS \ \forall p$, and we have all of the information to determine $D^S(t)/D^M(t)$, which in turn identifies γ^S/γ^M,meaning that I can identify the death rate up to a constant. In year 15, $D^S/D^m \cong 1.85$. I identify the final constant by setting the peak of my death curves equal to the empirical peak – this is the only point in the pictures which is set expressly to fit the data.

References

Ali, M. M., J. Cleland, and I. B. Shah (2004). Condom use within marriage: A neglected HIV intervention. *Bulletin of the World Health Organization 82*(3), 180–186.

Anderson, R. M., R. M. May, M. C. Boily, G. P. Garnett, and J. T. Rowley (1991). The spread of HIV-1 in africa: Sexual conact patterns and the predicted demographic impact of AIDS. *Nature 352*(6336), 581–589.

ASSA, the Actuarial Society of South Africa (2005). ASSA AIDS model 2002. www.assa.org.za. accessed February, 2005.

Beegle, K. and D. de Walque (2008). Demographic and socioeconomic patterns of HIV/AIDS prevalence in africa. In *The Changing HIV/AIDS Landscape: A Compendium of Selected Papers Prepared in Conjunction with the World Bank Africa Region's HIV/AIDS Agenda for Action 2007-2011*. forthcoming.

Bongaarts, J. (2006). Late marriage and the HIV epidemic in sub-saharan africa. Population Council Working Paper No. 216.

Castle, P. E., M. Schiffman, R. Herrero, A. Hildesheim, A. C. Rodriguez, M. C. Bratti, M. E. Sherman, S. Wacholdor, R. Tarone, and R. D. Burk (2005). A prospective study of age trends in cervical human papillomavirus acquisition and persistence in guanacaste, costa rica. *The Journal of Infectious Diseases 191*(1), 1808–16.

Center for Disease Control (2005). STD surveillance 2003. www.cdc.gov/std/stats/toc2003.htm. accessed August 27, 2005.

Central Bureau of Statistics (CBS) [Kenya], Ministry of Health (MOH) [Kenya], and ORC Macro (2004). Kenya demographic and health survey 2003. Calverton, Maryland: CBS, MOH, and ORC Macro.

Central Statistical Office (CSO) [Zimbabwe] and Macro International Inc. (2007). Zimbabwe demographic and health survey 2005-06. Calverton, Maryland: CSO and Macro International Inc.

Chen, C. Y., R. C. Ballard, C. M. Beck-Sague, Y. Dangor, F. Radabe, S. Schmid, J. B. Weiss, V. Tshabalala, G. Fehler, Y. Htun, and S. Morse (2000). Human immunodeficiency virus and genital ulcer disease in south africa: The herpetic connection. *Sexually Transmitted Diseases 27*(1), 21–29.

Clark, S. (2004). Early marriage and HIV risks in sub-saharan africa. *Studies in Family Planning 35*(3), 149–60.

de Walque, D. (2007). Sero-discordant couples in five african countries: Implications for prevention strategies. *Population and Development Review 33*(3), 501–523.

Downs, A. M. and I. D. Vincenzi (1996). Probability of heterosexual transmission of HIV: Relationship to the number of unprotected sexual contacts. *Journal of Acquired Immune Deficiency Syndromes and Human Retrovirology 11*(4), 388–395.

Dupas, P. (2007). Relative risks and the market for sex: Teenage pregnancy, HIV, and partner selection in kenya. mimeo, Dartmouth.

Fideli, U. S., S. A. Allen, R. Musonda, S. Trask, B. H. Hahn, H. Weiss, J. Mulenga, F. Kasolo, S. H. Vermund, and G. M. Aldrovandi (2001). Virologic and immunologic determinants of heterosexual transmission of human immunodeficiency virus type 1 in africa. *AIDS Research and Human Retroviruses 17*(10), 901–910.

Fleming, D. T. and J. N.Wasserheit (1999). From epidemiological synergy to public health policy and practice: The contribution of other sexually transmitted diseases to sexual transmission of HIV infection. *Sexually Transmitted Infections 75*(1), 3–17.

Gersovitz, M., H. G. Jacoby, F. S. Dedy, and A. G. Tape (1998). The balance of self-reported heterosexual activity in KAP surveys and the AIDS epidemic in africa. *Journal of the American Statistical Association 93*(443), 875–883.

Glynn, J., M. Carael, B. Auvert, M. Kahindo, J. Chege, R. Musonda, F. Kaona, A. Buve, and the Study Group on the Heterogeneity of HIV Epidemics in African Cities (2001). Why do young women have much higher prevalence of HIV than young men? a study in kisumu, kenya and ndola, zambia. *AIDS 15*(Suppl. 4), S51–S60.

Gray, R. H., M. J. Wawer, R. Brookmeyer, N. K. Sewankambo, D. Serwadda, F. Wabwire-Mangen, T. Lutalo, X. Li, T. VanCott, T. C. Quinn, and the Rakai Project Team (2001, April 14). Probability of HIV-1 transmission per coital act in monogamous, heterosexual, HIV-1-discordant couples in rakai, uganda. *Lancet 357*, 1149–53.

Heckathorn, D. (1997). Respondent-driven sampling: A new approach to the study of hidden populations. *Social Problems 44*(2), 174–199.

Jovanovic, B. (1979). Job matching and the theory of turnover. *The Journal of Political Economy 87*(5), 972–990.

Katzenstein, T. L. (2003). Molecular biological assessment methods and understanding the course of HIV infection. *APMIS 111*(Suppl. 114), 1–37.

Koopman, J. S., J. A. Jacquez, G. W. Welch, C. P. Simon, B. Foxman, S. M. Pollock, D. Barth-Jones, A. L. Adams, and K. Lange (1997). The role of primary HIV infection in the spread of HIV through populations. *Journal of Acquired Immune Deficiency Syndromes and Human Retroviruses 14*(3), 249–58.

Kremer, M. (1996). Integrating behavioral choice into epidemiological models of AIDS. *Quarterly Journal of Economics 111*(2), 549–73.

Kristof, N. D. (2005, March 30). When marriage kills. *The New York Times*. www.nytimes.com.

Lam, D., J. Seekings, and M. Sparks (2006, December). *The Cape Area Panel Study: Overview and Technical Documentation for Waves 1-2-3.* The University of Cape Town.

Laurian, Y., J. Peynet, and F. Verroust (1989). HIV infection in sexual partners of HIV-seropositive patients with hemophilia. *New England Journal of Medicine 320*(3), 183.

Lee, T.-H., N. Sakahara, E. Fiebig, M. P. Busch, T. R. O'Brien, and S. A. Herman (1996). Correlation of HIV-1 RNA levels in plasma and heterosexual transmission of HIV-1 from infected transfusion recipients. *Journal of Acquried Immune Deficiency Syndromes and Human Retroviruses 12*(4), 427–28.

Ministry of Health and Social Welfare (MOHSW) [Lesotho], Bureau of Statistics (BOS) [Lesotho], and ORC Macro (2005). Lesotho demographic and health survey 2004. Calverton, Maryland: MOH, BOS, and ORC Macro.

Morris, M. and M. Kretzschmar (1997). Concurrent partnerships and the spread of HIV. *AIDS 11*(5), 641–48.

National Center for Health Statistics (2002). National Survey of Family Growth, Cycle 6. http://www.cdc.gov/nchs/nsfg.htm.

National Statistical Office (NSO) [Malawi], and ORC Macro (2005). Malawi demographic and health survey 2004. Calverton, Maryland: NSO and ORC Macro.

Oster, E. (2005). Sexually transmitted infections, sexual behavior, and the HIV/AIDS epidemic. *Quarterly Journal of Economics 120*(2), 467–515.

Padian, N. S., S. C. Shiboski, and N. P. Jewell (1991). Female-to-male transmission of human immunodeficiency virus. *Journal of the American Medical Association 266*(12), 1664–67.

Pedraza, M.-A., J. Del Romero, F. Roldan, S. Garcia, M.-C. Ayerbe, A. R. Noriega, and

J. Alcami (1999). Heterosexual tranmsission of HIV-1 is associated with high plasma viral load levels and a positive viral isolation in the infected partner. *Journal of Acquired Immune Deficiency Syndromes 21*(2), 120–25.

Peterman, T. A., R. L. Stoneburner, J. R. Allen, H. W. Jaffe, and J. W. Curran (1988). Risk of human immunodeficiency virus transmission from heterosexual adults with transfusion-associated infections. *Journal of the American Medical Association 259*(1), 55–58.

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

Pictet, G., S. Le Coeur, P. M'Pele, N. Brouard, and M. Lallemant (1998). Contribution of AIDS to the general mortality in central africa: Evidence from a morgue-based study in brazzaville, congo. *AIDS 12*(16), 2217–23.

Pilcher, C. D., H. C. Tien, J. J. Eron Jr., P. L. Vernazza, P. W. Stewart, L.-E. Goh, and M. S. Cohen (2004, 15 May). Brief but efficient: Acute HIV infection and the sexual transmission of HIV. *Journal of Infectious Disease 189*, 1785–92.

Potts, M., D. T. Haperin, D. Kirby, A. Swidler, E. Marseille, J. D. Klausner, N. Hearst, R. G. Wamai, J. G. Kahn, and J. Walsh (2008). Reassessing HIV prevention. *Science 320*(5877), 749–750.

Quinn, T. C., M. J. Wawer, N. Sewankambo, D. Serwadda, C. Li, F. Wabwire-Mangen, M. O. Meehan, T. Lutalo, R. H. Gray, and the Rakai Project Study Group (2000). Viral load and heterosexual transmission of human immunodeficiency virus type. *New England Journal of Medicine 342*(13), 921–9.

Ragni, M. V., H. Faruki, and L. A. Kingsley (1998). Heterosexual HIV-1 transmission and viral load in hemophiliac patients. *Journal of Acquired Immune Disorders and Human Retroviruses 17*(1), 42–45.

Ragni, M. V., L. A. Kingsley, P. Nimorwicz, P. Gupta, and C. R. Rinaldo (1989). HIV

heterosexual transmission in hemophilia couples: Lack of relation to t4 numebr, clinical diagnosis, or duration of HIV exposure. *Journal of Acquired Immune Deficiency Syndromes 2*(6), 557–563.

Republic of South Africa Ministry of Health (2008). Progress report on declaration and committment on HIV and AIDS: Republic of south africa. In *United Nations General Assembly Special Session on HIV and AIDS*. www.unaids.org.

Saracco, A., M. Musicco, A. Nicolosi, G. Angarano, C. Arici, G. Gavazzeni, P. Costigiola, S. Gafa, C. Gervasoni, R. Luzzati, F. Piccinino, F. Puppo, B. Salassa, A. Sinicco, R. Stellini, U. Tirelli, G. Turbessi, G. M. Vigevani, G. Visco, R. Zerboni, and A. Lazzarin (1993). Man-to-woman sexual transmission of HIV: Longitudinal study of 343 steady partners of infected men. *Journal of Acquired Immune Deficiency Syndromes 6*(5), 497–502.

Sewankambo, N. K., R. H. Gray, S. Ahamad, D. Serwadda, F. Wabwire-Mangen, F. Nalugoda, N. Kiwanuka, T. Lutalo, G. Kigozi, C. Li, M. P. Meehan, H. Brahmbatt, and M. J. Wawer (2000). Mortality associated with HIV infection in rural rakai district, uganda. *AIDS 14*(15), 2391–400.

Smith, J. S. and N. J. Robinson (2002). Age-specific prevalence of infection with herpes simplex virus types 2 and 1: A global review. *The Journal of Infectious Diseases 186*(Suppl 1), S3–S28.

South Africa Department of Health and ORC Macro (2001). South African Demographic and Health Survey 1998. www.measuredhs.com.

Statistics South Africa (Producer) (2002-2004). South Africa Labour Force Survey, September 2002, September 2003, September 2004. Pretoria, South Africa: South African Data Archive (distributor). www.nrf.ac.za/sada.

Tanzania Government and ORC Macro (2005). Tanzania HIV/AIDS indicator survey 2003-04. Dar es Salaam: Tanzania Commission for AIDS and National Bureau of

Statistics.

Te Velde, E. R. and P. L. Pearson (2002). The variablility of female reproductive ageing. *Human Reproductive Update 8*(2), 141–54.

Templeton, A., J. K. Morris, and W. Parslow (1996). Factors that affect the outcome of in-vitro fertilization. *Lancet 348*(9039), 1402–06.

Van Noord-Zaadstra, B., C. Looman, H. Alsbach, J. Habbema, E. Te Velde, and J. Karbaat (1991). Delaying childbirth. effect of age on fecundity and outcome of pregnancy. *British Medical Journal 302*(6789), 1361–5.

Vincenzi, I. D. (1994). A longitudinal study of human immunodeficiency virus by heterosexual partners. *New England Journal of Medicine 331*(6), 341–46.

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica 57*(2), 307–333.

Wawer, M. J., R. H. Gray, N. K. Sewankambo, D. Serwadda, X. Li, O. Laeyendecker, N. Kiwanuka, G. Kigozi, M. Kiddugavu, T. Lutalo, F. Nalugoda, F. Wabwire-Mangen, M. P. Meehan, and T. C. Quinn (2005). Rates of HIV-1 infection per coital act by stage of HIV-1 infection, in rakai, uganda. *The Journal of Infectious Diseases 191*(9), 1403–9.

Wawer, M. J., N. K. Sewankambo, D. Serwadda, T. C. Quinn, L. A. Paxton, N. Kiwanuka, F. Wabwire-Mangen, C. Li, T. Lutalo, F. Nalugoda, C. A. Gaydos, L. H. Moulton, M. O. Meehan, S. Ahmed, and R. H. Gray (1999). Control of sexually transmitted diseases in uganda: A randomised community trial. *Lancet 353*(9152), 525–35.

Yerly, S., S. Vora, P. Rizzardi, J.-P. Chave, P. Vernazzi, M. Flepp, A. Telenti, M. Bettegay, A.-L. Veuthey, J.-P. Bru, M. Rickenbach, B. Hirschel, L. Perrin, and the Swiss HIV Cohort Study (2001). Acute HIV infection: Impact on the spread of HIV and transmission of drug resistance. *AIDS 15*(17), 2287–92.

age	% intercourse last week	% intercourse last 2 weeks	% intercourse last month
15-19	8.09	12.52	19.13
20-24	22.58	32.69	44.63
25-29	34.12	46.49	58.26
30-34	40.59	53.4	64.36
35-39	40.29	53.01	63.57
40-44	36.57	51.47	62.31
45-49	30.85	43.69	54.97

Table 1: Female Sexual Behavior by Age Group, South Africa DHS
Tabulations of sexual behavior by age, from the South African DHS

age	%birth in last year	Templeton adj %birth	van Noord-Zaadstra adj. %birth
15-19	7.08	N/A	11.91
20-24	15.3	15.3	15.39
25-29	15.3	15.3	15.57
30-34	12.56	13.59	14.98
35-39	8.91	11.99	20.09
40-44	2.63	7.69	11.48
45-49	0.41	2.79	N/A

Table 2: Births by age group (unadjusted and fecundity adjusted rates)
Tabulations of percent of births in last yeat by age, from the South African DHS. Column 2 adjusts births according to Templeton et al's (1996) estimates of fecundity, while column 3 utilizes van Noord-Zaadstra et al's (1991) estimates

Num Partners	NSFG Males	CAPS Males	NSFG Females	CAPS Females
	(1)	(2)	(3)	(4)
0	9.37	13.32	4.6	12.53
1	50.15	43.61	69.87	70.59
2	20.78	19.12	14.3	12.02
3	9.27	10.42	5.98	3.15
4	3.9	5.48	2.34	0.94
5	2.44	2.36	1.45	0.51
6	1.46	1.18	0.48	0.09
7+	2.63	4.5	0.96	0.18

Table 3: CAPS/NSFG Comparison: Tabulations of number of partners, last 12 months (sexually active adolescents only)
Tabulations of number of sexual partners in the last year. Columns 1 and 3 are using the NSFG data of US 14-22 year olds, while columns 2 and 4 utilize 14-22 year olds from Cape Town in the CAPS data

$\theta \in \{0, 1, 2, ..., 50\}$
$F(\theta) = N(25.5, 25/4)$
$\rho_1^m = \rho_1^f = .2$
$\rho_2^m = \rho_2^f = .02$
$\Pr(\theta'|\theta) = \left\{ \begin{array}{c} .3|\theta' = \theta \\ .2|\theta' = \theta \pm 1 \\ .15|\theta' = \theta \pm 2 \end{array} \right\}$
$T = 480$
$N = 100, K = 25$
$lifespan = 240$

Table 4: Parameter Inputs

var	median	mean	sd	min	max
lifetime partners	11	13.58	10.29	1	67
# partners, last 12 mos.	1	1.43	1.40	1	12
total years single	8.54	8.89	6.92	0.08	20

Table 5: Simulated Data Summary Statistics

δ	baseline		no acute infection	
	mean	sd	mean	sd
	0.072	0.006	0.039	0.003
0.01	0.062	0.005	0.039	0.003
	0.081	0.008	0.040	0.004
	0.136	0.009	0.076	0.004
0.02	0.121	0.008	0.075	0.004
	0.152	0.011	0.076	0.005
	0.189	0.008	0.110	0.004
0.03	0.170	0.008	0.111	0.005
	0.209	0.011	0.110	0.005
	0.239	0.009	0.141	0.004
0.04	0.216	0.008	0.141	0.006
	0.262	0.012	0.141	0.006
	0.282	0.009	0.173	0.005
0.05	0.257	0.008	0.174	0.006
	0.306	0.013	0.172	0.006
	0.321	0.007	0.203	0.006
0.06	0.294	0.008	0.203	0.007
	0.347	0.009	0.201	0.008

Table 6: Prevalence Results

Table presents simulated prevalences for each level of δ, the inputted prevalence from all other sources. In Columns 3 and 4, acute infection is shut down ($\rho_1^g = 0 \forall g$)

Yrs Married	Kenya	Lesotho	Malawi	Tanzania	Zimbabwe
5-9	0.032	0.019	0.056*	0.047***	0.110***
	(0.02)	(0.07)	(0.03)	(0.02)	(0.03)
10-14	0.076***	0.121	0.143***	0.052***	0.136***
	(0.02)	(0.08)	(0.02)	(0.02)	(0.03)
15+	0.100***	0.192*	0.121***	0.086***	0.235***
	(0.03)	(0.10)	(0.04)	(0.02)	(0.04)
N	897	558	928	1780	1427
Pseudo R-Sq	0.047	0.045	0.055	0.038	0.065
DepVar Mean	0.901	0.665	0.86	0.922	0.764
Year	2003	2004	2004	2003	2005/06

Table 7: HIV Neg. Couples by Marriage Tenure

Presents marginal effects from probits of a couple testing jointly negative for HIV on several categories of marriage tenure. Fixed effects for each 5-year age group and spousal age group are also included.

	Kenya		Lesotho		Malawi		Tanzania		Zimbabwe	
	Cluster	Individ	Cluster	Individ	Cluster	Individ	Cluster	Individ	Cluster	Individ
Age 1st Sex	-0.033***	-0.005**	-0.018	-0.010	-0.051***	-0.008**	-0.004	-0.002	-0.023***	-0.006*
	(0.01)	(0.00)	(0.02)	(0.01)	(0.01)	(0.00)	(0.01)	(0.00)	(0.01)	(0.00)
Age 1st Married	0.013*	0.007***	0.047***	0.023***	0.029**	0.009**	0.006**	0.004***	0.021***	0.016***
	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)		(0.00)	(0.00)	(0.01)	(0.00)
Polygamous	0.023				-0.218***		0.012		-0.049	
	(0.06)				(0.07)		(0.05)		(0.05)	
Spouse Age Gap	-0.002	0.001	0.004	0.001	0.013**	-0.004	0.000	0.000	0.006**	0.004***
	(0.00)	(0.00)	(0.01)	(0.00)	(0.01)	0.003*	(0.00)	(0.00)	(0.00)	(0.00)
Mean Prev		0.203***		0.069		0.253***		0.322***		0.151**
		(0.04)		(0.06)		(0.04)		(0.04)		(0.06)
N	394	1415	403	1425	499	1256	345	2824	398	3045
R-sq (adj. or ps)	0.037	0.064	0.051	0.015	0.047	0.060	0.000	0.084	0.042	0.022

Table 8: HIV and Duration of Singlehood

"Cluster" columns present ols regressions of Sampling Cluster prevalence on population mean characteristics. "Individ" Columns present marginal effects on infection from an individual-level probits, where mean prev is the mean HIV prevalence among other individuals in the sampling cluster.

Years Single	Kenya	Lesotho	Malawi	Tanzania	Zimbabwe
Years Single*	0.006**	0.035***	0.009*	0.007***	0.028***
Married 0-4 yrs	(0.002)	(0.009)	(0.005)	(0.002)	(0.005)
Years Single *	0.009***	0.012	0.017**	0.007***	0.013***
Married 5-9 yrs	(0.003)	(0.014)	(0.007)	(0.002)	(0.005)
Years Single *	0.004	0.018	0.006	0.000	0.012*
Married 10-14 yrs	(0.006)	(0.012)	(0.009)	(0.004)	(0.006)
Years Single *	0.004	-0.017	0.004	0.004	-0.002
Married 15+ yrs	(0.004)	(0.020)	(0.011)	(0.004)	(0.008)
Observations	1415	1425	1256	2824	3045
Pseudo R2	0.043	0.019	0.017	0.023	0.027
Mean HIV	0.069	0.259	0.106	0.046	0.164

Table 9: Years Single on HIV by Marital Tenure

Presents marginal effects from a probit of HIV status on marital tenure for women. Dummies for each marital tenure group are also included as covariates

λ	Mean Prev	Safe Prev	Unsafe Prev
0	0.189		0.189
	(0.008)		(0.008)
0.05	0.182	0.141	0.185
	(0.009)	(0.014)	(0.009)
0.1	0.178	0.141	0.183
	(0.008)	(0.010)	(0.008)
0.25	0.158	0.131	0.168
	(0.006)	(0.006)	(0.007)
0.5	0.130	0.115	0.146
	(0.007)	(0.007)	(0.007)
0.75	0.108	0.101	0.128
	(0.004)	(0.004)	(0.007)
1	0.092	0.092	
	(0.004)	(0.004)	

Table 10: Policy: Condoms in Short Relationships

Presents lifetime prevalences when $\delta = .03$ and fraction λ use condoms in short relationships. "Mean Prev" gives the overall rate, "Safe Prev" gives the rate for people who wear condoms in short relationships, and "Unsafe Prev" gives the prevalence for those who do not. Standard errors in parentheses.

Study	Location	Year	Report	N (M/F)	Per Partner	PPY	PPY-at risk	Per Coital Act	Sex per mo
Gray et al	Uganda	2001	PCA	174	21.84%	11.60%	11.60%	0.0011	8.92
M/F Comparison				97/77	17.5%/27.3%	9.1%/14.8%	9.1%/14.8%	0.0009/0.0013	8.27/9.74
Fideli et al	Zambia	2001	PPY	1022	1022	15.8%	7.70%		
M/F Comparison				535/487	12.33%/8.8%	8.28%/7.07%			
Quinn et al	Uganda	2000	PPY	415	21.70%	11.80%	11.80%		
M/F Comparison				228/187	21.4%/21.9%	12.0%/11.6%	12.0%/11.6%		
Pedraza et al	Spain	1999	PP	38	26.30%	5.76%	5.76%	0.0005	9
M/F Comparison				27/11	25.9%/27.3%				
Downs and De Vincenzi	Europe	1996	PP	121	9.90%	5.00%	5.00%	0.0004	3.5
M/F Comparison				73/48	11.0%/8.3%			0.0005/0.0003	
Rockstroh et al	USA	1995	PP	198 (M)	0.100				
De Vincenzi	Europe	1994	PPY	245	4.70%	2.30%	4.80%	0.001	
M/F Comparison				157/89	5.09%/4.6%				
Saracco et al	Europe	1993	PPY	343 (M)	0.055	3.90%	7.20%		
Padian et al	USA	1991	PP	379	16.40%				
M/F Comparison				307/72	19.9%/1.4%				
Ragni et al	USA	1989	PP	45 (M)	0.130				
Laurian et al	Europe	1989	PP	31	0.097	0.0480	0.0880		
Peterman et al	USA	1988	PP	80	15%			0.0099	4.36
M/F Comparison				55/25	18.2%/8%			0.0012/0.0005	4.29/4.52

Table 11: Review of Medical Studies

Report indicates the chosen method of reporting, where PCA = per coital act, PP = partnership, PPY = per partnership-year. At-risk partnerships are defined to be those who do not report perfect condom use.

$Tratio$	λ	Mean Prev	Safe Prev	Unsafe Prev
0.875	0	0.152		0.152
		(0.008)		(0.008)
0.75	0	0.120		0.120
		(0.006)		(0.006)
0.75	0.5	0.093	0.084	0.103
		(0.004)	(0.004)	(0.006)
0.75	1	0.074	0.074	
		(0.003)	(0.003)	

Table 12: Changing transmission rates

Presents evidence on lifetime prevalence when $\delta = 0.03$, transmission rates are multiplied by Tratio, and fraction λ use condoms in short relationships."Mean Prev" gives the overall rate, "Safe Prev" gives the rate for people who wear condoms in short relationships, and "Unsafe Prev" gives the prevalence for those who do not. Standard errors in parentheses.

	Kenya	Lesotho	Malawi	Tanzania	Zimbabwe	South Africa
Individual Sum Stats						
HIV	0.073	0.259	0.106	0.046	0.164	
	(0.008)	(0.012)	(0.009)	(0.004)	(0.007)	
Age	30.202	30.801	27.658	29.232	29.294	34.270
	(0.224)	(0.239)	(0.236)	(0.151)	(0.153)	(0.144)
Age First Sex	17.074	16.652	16.392	17.183	17.930	16.908
	(0.085)	(0.069)	(0.068)	(0.053)	(0.051)	(0.042)
Age 1st Marriage	19.155	18.358	17.338	18.371	18.705	21.287
	(0.104)	(0.080)	(0.077)	(0.066)	(0.060)	(0.970)
Spousal Age Gap	6.379	5.986	5.508	6.826	6.367	6.153
	(0.129)	(0.135)	(0.116)	(0.113)	(0.094)	(0.103)
N	1415	1425	1256	2824	3045	2915
Polygynous (in cluster)	0.162		0.160	0.092	0.107	0.048
	(0.008)		(0.006)	(0.005)	(0.006)	(0.004)
Cluster N	394	402	499	345	398	935
Couple Sum Stats						
Both Negative	0.901	0.665	0.860	0.923	0.764	
	(0.010)	(0.020)	(0.011)	(0.006)	(0.011)	
Years Married	10.217	11.857	9.680	8.996	10.227	
	(0.271)	(0.400)	(0.259)	(0.157)	(0.214)	
Couples N	897	558	928	1780	1427	
Year	2003	2004	2004	2003	2005/06	1998

Table 13: Means from DHS data

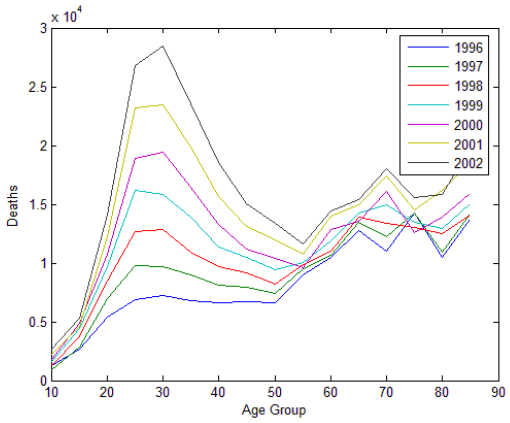

Figure 1: Women's Deaths by Age, 1996-2002

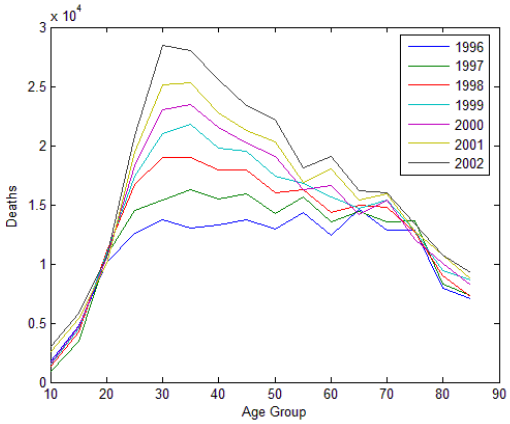

Figure 2: Men's Deaths by Age, 1996-2002

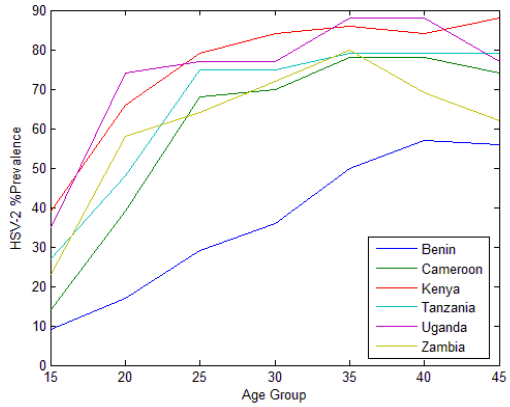

Figure 3: Female HSV-2 Prevalence by Age

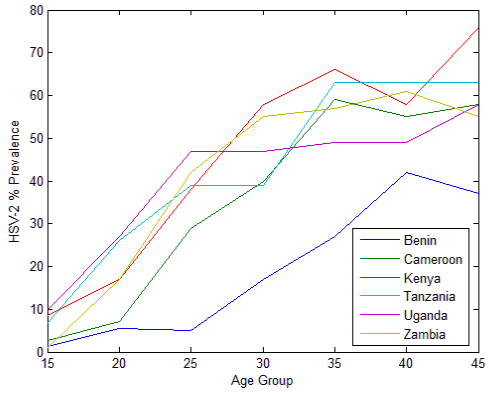

Figure 4: Male HSV-2 Prevalence by Age

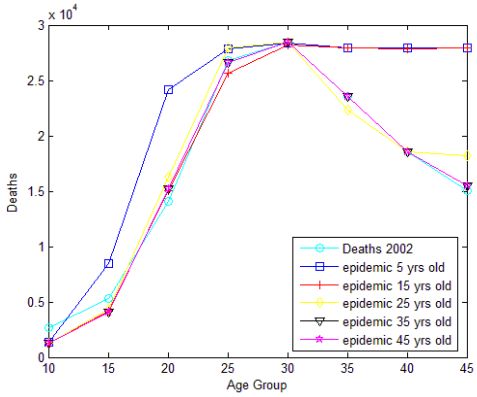

Figure 5: Best Fits at Various Epidemic Ages

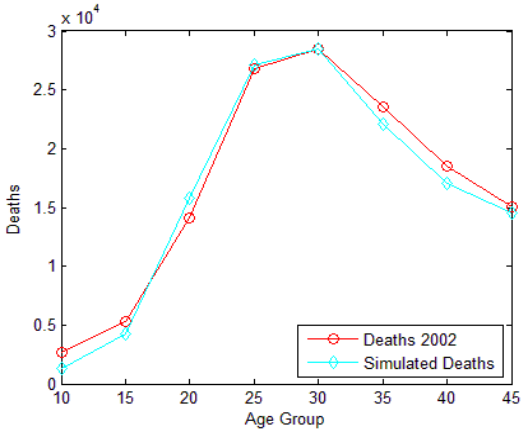

Figure 6: Women's Deaths in 2002

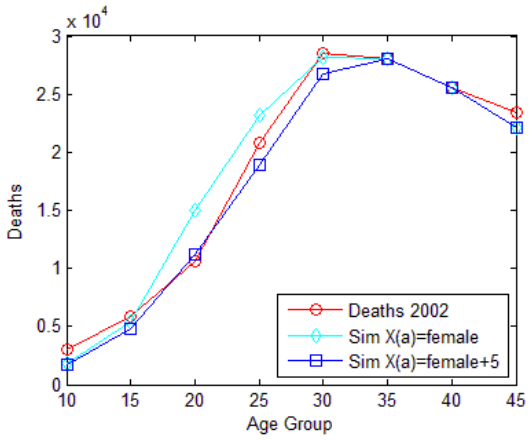

Figure 7: Men's Deaths in 2002

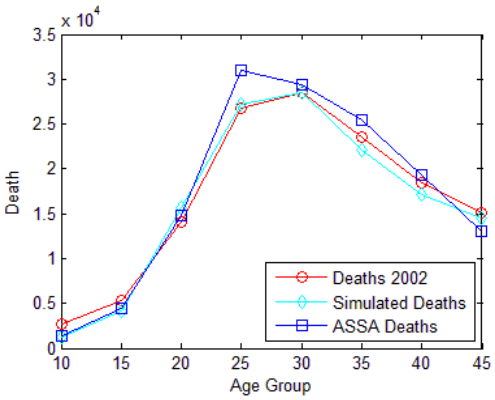

Figure 8: 2002 Actual, Predicted, and ASSA Women

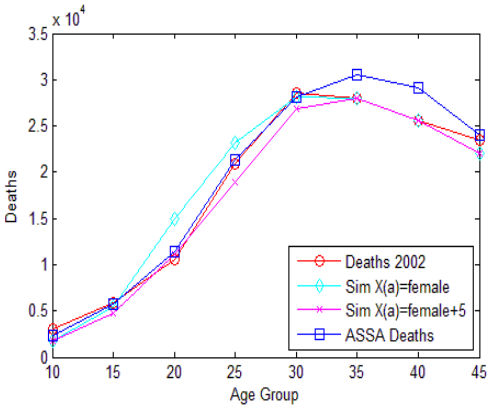

Figure 9: 2002 Actual, Predicted, and ASSA Men

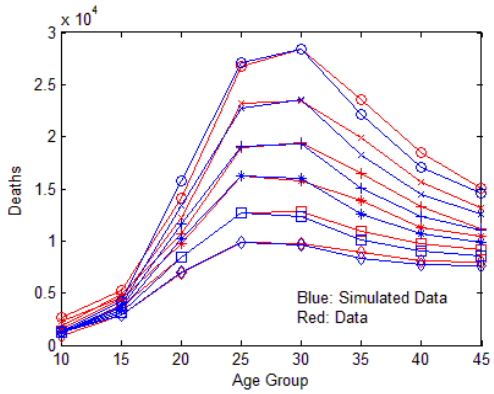

Figure 10: Women's Deaths, 1997-2002

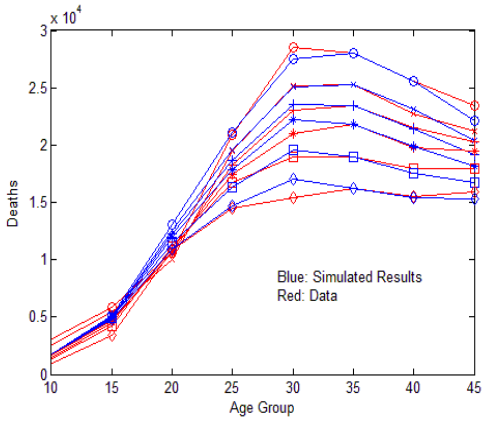

Figure 11: Men's Deaths, 1997-2002

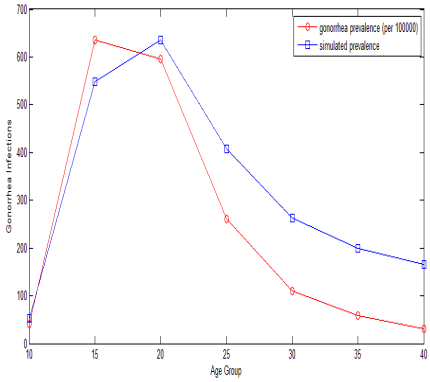

Figure 12: US Female Gonorrhea prevalence by Age

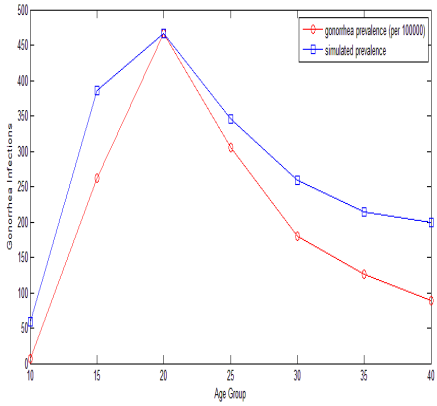

Figure 13: US Male Gonorrhea by Age