

# Marital Shopping and Epidemic AIDS

Jeremy R. Magruder\*

January 29, 2008

## Abstract

The high prevalence rates of HIV in much of Africa seem especially striking when considered alongside the very consistently estimated, very low transmission rate of the virus. Death registries in South Africa reveal a strong age-trend in infection risk – nearly all women who will ever be infected are infected by age 30 or so, a decade younger than any measurable declines in coital frequency. This paper argues that existing behavioral epidemiological models based on heterogeneous risk are incapable of generating this age-trend, particularly in the early epidemic phase which is relevant to South Africa. Adapting Jovanovic’s (1979) model of job turnover into a model of relationship turnover, I show that searching for a partner with a good match quality generates exactly the timing needed to turn a very brief window of high infectiousness described in the medical literature into high prevalence rates – in other words, simple dating behavior is in itself sufficient to create an epidemic. Using DHS data, I construct the age-profile of deaths suggested by matching models, and find that it fits the actual deaths data very well. Model predictions are compared to reported sexual behavior in South Africa, HIV prevalence in Kenya, and gonorrhea in the US, finding a good fit.

## 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 al 2001), or 8-12% per partner-year. Despite this, antenatal prevalences reach as high as 40% in Botswana and 25% in South Africa, suggesting

---

\*University of California, Berkeley. e-mail [jmagruder@berkeley.edu](mailto:jmagruder@berkeley.edu). I am indebted to Fabian Lange, Nicoli Nattrass, Rohini Pande, Paul Schultz, Jeremy Seekings, Chris Udry and seminar participants at Yale, Berkeley, and the University of Cape Town for many helpful suggestions. All mistakes are naturally my own.

that a large fraction of the population have been afflicted by this extremely unlikely event. 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 appears to occur 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<sup>1</sup>. 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 25 for women (which is notably the median age of first marriage for

---

<sup>1</sup>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.

South African women). This paper is the first to show that a simple search and matching model of serial monogamy is 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 sexual behaviors. Using DHS data on marital behavior and theoretical predictions of search-and-matching matching models, I show that matching behavior generates an age-death profile which fits that observed in South Africa extremely well in a way robust to a large variety of parametric assumptions. I further show that the length of singlehood – a key predictor of HIV in a matching model – is strongly correlated with HIV infection in Kenya in a way that other classic predictors (e.g. spousal age gaps and polygamy) are not, and that this correlation only exists for the recently single, so that it is not merely serving as a proxy for risk preferences. I show that reported sexual behavior more closely resembles matching behavior than heterogeneous risk preferences. Given that reported behavior looks similar in the US and Cape Town, I examine whether the marital shopping model can predict the age-profile of Gonorrhoea in the United States and find a good fit. Finally, I evaluate the public health implications of Marital Shopping as an important vector for the spread of HIV.

## **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 observed in South Africa, where death registration is relatively complete. Like other parts of Southern Africa, epidemic AIDS is a recent phenomenon in South Africa. 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<sup>2</sup>.

As the focus of this paper is on assessing the ability of various models to 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. A 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, a biomarker for sexual risk, namely pregnancy, is preferred. Table 2 reports pregnancy rates by age. Indeed, pregnancy rates do decline with

---

<sup>2</sup>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.

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 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)<sup>3</sup>. 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.

### 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

---

<sup>3</sup>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.

(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 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<sup>4</sup>.

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<sup>5</sup>. As reported in columns 2 and 4 of table 3,

---

<sup>4</sup>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, 1960 then such a network could exist, but not if everyone matches with people who are within one or two years of their age.

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

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 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. The second assumption is necessary as this age-independent explanation revolves around the high-risk older women having already attrited from the population, which seems unlikely at the onset of the epidemic (and, given how recently there were minimal deaths from HIV in South Africa, it is safe to assume that the epidemic is in its early years). To discover whether an age-independent story can be successful once we relax these two assumptions which seem unlikely in this context, I calibrate a heterogeneous risk model with three types.

#### **4.1 Age-Independent Model Calibration**

I follow two strategies to show that both stability of the epidemic and a relatively extreme assumption on the distribution of risk are necessary for heterogeneous risk to predict the South African age-death profile. In the first, I assume stability in the HIV epidemic but require there to be some fraction of individuals where infection is not deterministic. While we can generally imagine heterogeneous risk levels to be distributed in any way, I restrict attention to the case where there are three types: one who is at high risk, one at low risk, and one at negligible risk. In this calibration, I show that if there exists one type of individuals who represent a reasonable fraction of the at-risk population and are at palpable but not deterministic risk, then heterogeneous preferences cannot create an age-death profile which closely resembles the South African case, even in the steady state. I presume that these groups were each subjected to their type's risk level each year that they remained uninfected, assuming that incidence rates have been constant throughout all age-groups' lifetimes. This is clearly an unrealistic assumption for South Africa in 2002 where the epidemic was a recent phenomenon and prevalence rates had been growing quickly; however, it avoids the difficulties of assumptions about the timing and speed with which the epidemic

struck South Africa. It also is a “worst-case” scenario for the rejection of heterogeneous, static preferences as an underlying source of the epidemic, as I illustrate when I relax this assumption below. Then, if  $\rho_g$  is the annual infection risk of group  $g$ ,  $f_g$  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, the number of deaths at age  $a$  would be

$$D(a) = POP * \sum_{g=1}^2 \sum_{k=0}^{a-1} \sum_{\tau=0}^{a-k} (X(k) - X(k+1)) * f_g (1 - \rho_g)^{(\tau-1)} \rho_g (S(a - \tau) - S(a - \tau + 1)) \quad (1)$$

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. This parameter is of independent interest, as it indicates the number of individuals who are at any risk whatsoever. We also know from census data the number of women living in South Africa at this time period. Therefore, the difference between this estimated number and the population represents a third group, who are at zero risk. I then propose several constraints: in table XXY, tables 1 through 3 restrict  $\rho_1$  to be less or equal to .2, .1, and .05 respectively (.2 is the highest considered as this is the annual risk from unprotected sexual activity with a spouse who is infected among young couples who have high coital frequencies (Gray et al 2001)). There are three other free parameters in this model,  $f_1$ ,  $\rho_1$  and  $\rho_2$ . In figures 5 through 7, I show the best possible fits with various constraints on these parameters. Restrictions on  $f_1$  force the model to allocate various fractions of individuals to group 2; in order to ensure that group 2 is indeed different from group 1, I restrict  $\rho_2$  to be .02<sup>6</sup>. At this rate, about 55% of individuals will be infected within the 40 years considered, so while this does indeed represent a high risk level,

---

<sup>6</sup>Without this restriction, the data often prefers to set  $\rho_2 = \rho_1$  or  $\rho_2 = 0$ , returning us to the two-type case. Of course, the choice of .02 is a bit ad hoc. Similar results are achieved for other low but real risk levels of  $\rho_2$ , such that as the  $\rho_2$  approaches zero the maximum size of  $f_1$  which generates a similarly bad fit also approaches 0. With  $\rho_1 \leq .1$ ,  $f_1 \leq .2$ , the best possible fit is worse than the marital shopping model for  $\rho_2 \geq .005$ .

it is chosen to be a "middle risk level" such that some individuals will eventually be infected while others may not. Figures 5 through 7 reveal that, while it is possible to mimic the age-death distribution quite closely when as much as half of the at-risk population is at very high risk (10% or 20% per year), it mimics the distribution less well as the more and more of the at-risk population is lower risk (and, indeed, less well than the marital shopping model) or when high-risk individuals are at 5% annual risk – still a level which leads to almost deterministic infection over a lifetime (this level is achieved if individuals have unprotected intercourse as frequently as the young adults in the Gray et al (2001) study, but are only able to find an infected partner 25% of the time). Moreover, in order to generate these death numbers with these risk rates, the vast majority of the population must be at zero risk (or, at extremely low risk) – in all of the figures presented, the total at-risk population (*POP*) is no more than about 2% of the population of 15-50 year old women in South Africa. We've already established that a large fraction of at-risk individuals must be at high-risk, and nearly all of these high risk individuals will become infected, so since there are relatively few deaths by 2002 compared to the size of the population, it must be the case that most individuals are safe. In the three type case, in order to capture this sharp peak, it must be the case that the vast majority are at very low risk levels, while a large fraction of those at risk are at very high risk. This scenario does not accord very well with the data in table 3, if risk levels are correlated with preferences for partner variety.

However, the few deaths relative to the population in 2002 brings about another point. Relatively few individuals were dying annually compared to the population size in 2004 in part because the epidemic was very young – indeed, this paper has argued that the number of deaths in 1996 was a good approximation for the deaths by age when no AIDS deaths were taken into account. The above analysis was performed assuming a stable age-death distribution. If we change the risk calculations to allow for the fact that the beginning of reasonable risk levels for South Africans could not have been more than 15 years or so ago (as there are so few AIDS deaths 7 years thereafter), this changes the pictures generated sharply.

Unfortunately, little data exists to accurately describe the dynamics of the incidence rate before prevalence was as high as it is today. To give preferences the best possible chance at predicting this age-death curve a few years into the epidemic, I search over several parameters to find the best fit. This suggests a slightly different calculation. Now,

$$D(a, 2002) = POP * \sum_{g=1}^2 \sum_{k=0}^{a-1} \sum_{\tau=0}^{a-k} (X(k) - X(k+1)) f_g * \prod_{y=1987}^{2002-\tau} (1 - \rho_g(y)) \rho_g(\tau) (S(a-\tau) - S(a-\tau+1)) \quad (2)$$

As above,  $\rho_1$  is allowed to be between 0 and .2 (at it's peak), and I assume that it increases linearly until it reaches the peak incidence rate. Since I do not know when the epidemic would have reached peak incidence years 1 through 15, I find the best fit allowing any of these years as the first year of peak incidence. As it turns out, at all phases of the epidemic, year 1 is chosen as the best fit, meaning that the age-death chart most resembles the unrealistic situation where peak incidence was achieved immediately. For this case, I relax the restrictions on  $f_1$  and  $\rho_2$ , so that the underlying preference parameters can be any distribution of one, two, or three types<sup>7</sup>. Therefore, any fraction up to 100% of the at-risk population is allowed to be in the high-risk group – as shown above, this can generate a very good fit with stable risk levels. However, when the infection is young, as figure 8 illustrates, such is not the case. In fact, around year 15, the peak infectiousness is amongst the oldest groups considered, and the close fit to the age-death curve is not achieved until 35 years or so after peak infectiousness began. Preferences 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, so that many risky people were safe in their youths, or when many people prefer low or middle risk levels, it can only provide a poor approximation.

---

<sup>7</sup>Here,  $\rho_2$  is allowed to be between 0 and 0.02, and increases linearly at the same rate that  $\rho_1$  does.

## 5 Marriage as a risk factor

Given that age-independent models of heterogeneous risk are unable to generate the age-profile of infection 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 common story about the implications of marriage for HIV, which describes how (previously low-risk) wives are now subjected to the continued risk that their husbands assume. For example, after New York Times Columnist 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 infection likelihood, we can rule out the cheating husbands story as being the mechanism by which most people become infected.

The discussion of the dangers of marriage in Africa is so pervasive that, unsurprisingly, there is some empirical work. For example, Clark(2004) finds that young women (aged 15-19) who are married are much more likely to test HIV-positive than those who are not, though the point estimate loses significance after conditioning on age and other covariates. Examining reports of sexual behavior, she concludes that early marriage increases risky behavior among young women and puts them at greater risk. Another interpretation would be that unmarried individuals differ from married individuals in that they never have found a marriage-worthy partner; in particular, they may differ as to the level of effort that they’ve put into such a search. In this case, those who undergo this process early differ only in the age at which they expose themselves to this risk from those who delay marriage. In fact, evidence supporting this latter interpretation is found by Bongaarts (2006) who discovers

that married individuals have higher infection rates, but that this result is due to the fact that they have been sexually active longer – a year spent single and sexually active is riskier than a year spent married. Glynn et al (2001) similarly find supportive evidence, as they observe that lifetime prevalences remain much lower among women who report being virgins at marriage, inspiring the suggestion that “much of the HIV in women is acquired before marriage.” Getting married appears to be dangerous – but staying married may be relatively safe. Indeed, if spousal search is an important source of HIV infection, then this would indeed suggest age-declining risks, as individuals undergoing spousal search are systematically younger than those who are married. However, it is not yet established either whether behavior during spousal search is capable of generating epidemic HIV prevalence rates, nor if the particular age-risk curve predicted by spousal search would resemble that in South Africa. The next sections of this paper confirm these two facts.

## 6 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  $\theta_{ij}$  is 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. 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

$$V_t(\theta_{ij}) = \max_{\{stay, leave\}} \left\{ \theta_{ij} + q_i + \beta E[V_{t+1}(\theta'_{ij}) | \theta_{ij}], E[\theta_{ij'} + q_i + \beta V_{t+1}(\theta'_{ij'})] \right\} \quad (3)$$

where  $\beta$  is the discount rate. 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. 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 median 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 where the initial signal is much more informative than subsequent ones. 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<sup>8</sup>.

## 6.1 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 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, or with survey follow-ups, per contact. The studies in the Rakai district, Uganda, are particularly compelling as participants report minimal condom use despite counselling. Here, average

---

<sup>8</sup>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 (this difference is much finer than the discretization used in numerical calculations). 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.

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 high condom usage was reported, was observed to be somewhat less (8% PPY) (Fideli et al 2001). In fact, these rates (8-12%) aren't very much different from those observed in the US and Europe (5-10% PPY) despite lower reported condom use, higher reported coital frequencies, and the much higher burden of other sexually transmitted infections, inspiring Gray to note "greater infectivity of predominant HIV-1 viral subtypes is unlikely to account for the explosive HIV-1 epidemic in sub-Saharan Africa" (2001).

However, the average infection rate may not be a sufficient statistic to understand the dynamics of HIV transmission. Indeed, while per-partnership-year transmission rates are useful for understanding average transmission patterns, they 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. Moreover, as viral load in genital secretions has been shown to be correlated with blood viral load but not perfectly aligned (Pilcher et al 2004) with what looks something like classical measurement error, these estimates may be subject to attenuation bias, meaning viral load may be even more important than the authors conclude. 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 soars to huge levels. 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)<sup>9</sup>. Koopman et al (1997) generate an epidemiological model assuming individuals transition randomly between high and low-turnover states and match non-randomly with others in the same state. Considering the US, and using larger transmission probabilities associated with homosexual behavior, 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  $\rho_1^g$  ( $\rho_2^g$ ) 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 ( $\delta$ ) enter already infected with HIV, with  $\delta/6$  entering with equal probability in the first, second, or third month of acute infection<sup>10</sup>. This consistent injection of HIV allows me to obtain higher prevalences and can be interpreted

---

<sup>9</sup>These dynamics of infections have only recently become well-understood. Prior to this understanding, some epidemiological researchers (e.g. Downs and Di Vicenzi 1996) proposed a per-partnership transmission rate rather than a per-contact or duration-dependent rate, as heterogeneity in infectiousness led to an infection profile which was not binomial in sexual contacts. This approach has been adopted in much of the economics literature (e.g. Kremer 1996, Oster 2005). This paper follows current epidemiological practice in taking advantage of these more recent advances in medical understanding to explicitly model these dynamics.

<sup>10</sup>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.

as infection from all other sources. My preferred interpretation for this injection is that it represents new matches for individuals whose spouses may have died or migrants who find new partners; however, all other hypothesized sources of HIV infection are included within  $\delta$  as well (including intravenous drug use, men who have sex with men, prostitution, infidelity, concurrency, etc.). Individuals who are initially infected do not behave differently from other individuals as high risk individuals would in a “core group” model.

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 are follow Pilcher et al (2004) and Wawer et al (2005). The distribution over  $\theta$  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 my simulated women.

To simulate this model, I first solve the dynamic programming problem numerically according to the distributional assumptions given in table 4. Then, each man and women in the first cohort randomly draws a partner and receives a draw of partner quality. If their partner is HIV positive, they face a random draw on the probability of contracting HIV themselves. At the end of the period, they decide whether to stay in the relationship or draw a new one based on whether or not it exceeds their reservation quality at that point in their lives. If they stay, their relationship evolves and they face the same decision in the subsequent period. If they leave, they receive a new partner and a new draw on theta. Every four months, a new cohort of  $N$  individuals enters the pool of searching singles, and I simulate the model for 50 years. Table 5 features summary statistics from the simulated data.

## 7 Results

Table 6 reports the fraction of women who will become infected for each  $\delta$  (summary statistics are means of 50 simulations). What I find, in practice, is that every percentage point of infected individuals which I input into the model results in about 7 or 8 additional infections. In other words, each person who enters the spousal search period infected will cause about three other single people to become infected, all of whom 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 2-3% 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 will reach around 80% of the level that it ever will. By the end of the epidemic, each infection is responsible for about 8-9 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 reasons for this is the basic behavior which is generated by matching models (and hence fairly insensitive to parametric assumptions). Individuals have clusters of several very short relationships which are 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, and disproportionately likely to remain in the acute phase of infection and pass their infection on to 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. Table 7 reports the prevalence results of this situation. The results are strong: three months of protection cuts the epidemic prevalence rate in half. Now,  $\delta$  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 which has shown that a large number of infections are attributable to acute infection. Yerly et al (2001), in a study in Switzerland, perform very specific analysis on the RNA of different individuals' HIV in Switzerland and is able to trace clusters of infections. Dealing with under a tenth of estimated new infections over four years, they are still able to trace a third of them into specific clusters of outbreaks, where individuals were infected with very similar strands of the virus in the same few months in the same area. As they have only a small percentage of total infections, they can't observe some clusters, so attributing a third of the infections to acute infection-based outbreaks is an underestimate. 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 familiar and intuitive; it is very similar to public health messages which were given out in response to 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).

Does the matching model can 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 when single and searching for the uninfected. 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 transmission rate times the risk that your spouse entered the marriage already infected times the risk that you entered the marriage uninfected times the probability that you haven't caught it from your spouse yet. 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<sup>11</sup>. 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, on average, five years later than women. As described above, I know the risk profile 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. This allows me to predict a age-profile of infection for the single and for married people. Therefore, to construct the age-death profile, I need time paths of the prevalence for single people, for newlyweds, and the relative fractions of infections which take place when single and when married. Each of these are taken from simulations; reassuringly, the figures below are robust to reasonable changes in all of them as it is the empirical distributions of singlehood and marriage as well as the theoretical risk predictions which fundamentally underly the close fit.

Figures 9 and 10 represent this comparison. Here, I make the further assumption that the

---

<sup>11</sup>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 (and in some cases superior) fits.

number of deaths by age in 1996 represents the number of non-AIDS deaths in 2002, and I set the peak of my AIDS deaths equal to the empirical peak. That is, the 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 underly this difference; however, the important point is that the quantity of infections that this non-monogamy is responsible for is tiny in comparison to the monogamy. For men, the empirical data appears to lie in between the age death curves created by assuming that men have the same sort of sexual initiation as women and that assuming that their age-pattern in sexual initiation is shifted 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 above, is the most successful a heterogeneous risks model can achieve (given in figure 8). 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, give 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.<sup>12</sup> Figures 11 and 12 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 sums of squared deviations of about half the magnitude.

---

<sup>12</sup>Notably, marriage is not considered in the ASSA model.

## 8 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.

One 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 lifetime match, and should be correlated with the amount of risk assessed over a lifetime. Bongaarts (2006) discusses cross-country and individual level data from DHS surveys which suggest that both age at marriage and the gap between the age at first sex and age at marriage are important predictors of HIV infection or prevalence. Similar to Bongaarts's analysis, table 8 predicts HIV prevalence at the sample cluster and individual level using the 2003 Kenya DHS survey. Here, I estimate

$$\bar{Y}_i = \beta_1 (\overline{Age1stSex}_i) + \beta_2 (\overline{AgeMarried}_i) + \beta_3 (\overline{Polygamy}_i) + \beta_4 (\overline{AgeDiff}_i) + \varepsilon_i$$

Where  $\bar{Y}_i$  is the HIV prevalence in sampling cluster  $i$ ,  $\overline{AgeDiff}_i$  is the difference in age between spouses, and variables represent sampling cluster-level means. If the length of singlehood is the important variable, then we should observe that the age of 1st sex and age of first marriage to be opposite in sign and similar in value. Column 1 and 2 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. These results are made more striking by the fact that two other likely candidates for HIV risk, polygamy and spousal age differences, are not statistically significant, though the polygamy coefficient is large and noisy. One concern may be that sex and marriage behaviors have adjusted to accommodate the differential HIV

prevalences in sampling clusters. To correct this, Columns 3 and 4 repeat the analysis using sex and marriage variables for women over 40, for whom these decisions will have been made years earlier, before there were many HIV cases. Once again, singlehood seems to be the primary determining variable.

A natural extension is to examine whether the correlation is present at the individual level as well as the sampling cluster. Indeed, Columns 5 and 6 reveal that this same characteristic is also correlated at the individual level, conditional on cluster HIV prevalence. 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. Fortunately, a simple test is available: if singlehood is risky when single, then the effect of the married-single gap should be strongest over the recently married, as those who have been married longer would have been single when HIV was less prevalent and are likely to be absent from the sample due to premature death if they caught HIV when Single. In contrast, if a long period of singlehood is simply correlated with other risky behaviors, such as preferences, then those who were single for a long time period should still be at risk years after marriage. Figure 15 presents marginal effects from a probit of HIV prevalence on average cluster HIV prevalence and years of singlehood, where the effect of singlehood on HIV and the average HIV prevalence are allowed to change by years of marriage. As is apparent, a long singlehood is quite risky for the first five years of marriage, and declines slightly in point estimate and remains different from zero for the second half-decade. After longer marriage tenures, the point estimate becomes quite small and loses its significance (though it becomes noisier as well). 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. Moreover, we can reject the hypothesis that the effect of being single is the same for those married within the last five years to those married more than 15 years ago at the five percent

level.

Another prediction of the marital shopping model is that, when single, people should have a large variance in their annual number of sexual partners. In particular, a year of singlehood where you have a sequence of bad draws in partner quality will result in a large number of partners, whereas you may have very 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 individual 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 model by simply multiplying the annual number of partners with the years of sexual activity, subject only to integer problems<sup>13</sup>.

Tables 9 through 12 show "kitchen sink" regressions of sexual behavior variables which seem likely to be correlated with matching or Kremer's model for men and women, along with results from two sources. Tables 9 and 10 utilize the Cape Area Panel Study, a panel dataset of young adults aged 14-22 in 2002 who live in the Cape Town Metropolitan Area, for males and females respectively. In 2002, the number of partners from that year was asked, while in 2003 a random subsample was reinterviewed and lifetime number of partners was asked. Very few of these young adults have become married, so although it seems important to the matching model I exclude it from the model selection test for these young adults. Unfortunately, sample size is small, which creates problems of statistical significance in the OLS estimates. In order to create a comparable sample from my simulated data, I use simulated agents with an identical distribution of years of sexual activity as that reported in the CAPS data. Tables 11 and 12 utilize the National Survey of Family Growth (NSFG), an American dataset of men and women aged 15-44 for males and females. Tobits are used,

---

<sup>13</sup>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.

as for men annual and lifetime partners are right-truncated at 7, and for women lifetime partners are right-truncated at 50. I exclude observations where the relationship between annual and lifetime partners is deterministic, that is, observations where the individual has been sexually active for a year or less or men who have had seven or more partners in the previous year. Column 1 of Tables 9 through 12 report the estimates from the data, while column 2 reports the results from running identical estimation on simulated data from the matching model and column 3 illustrates the theoretical predictions of Kremer’s model. For men and women in the US and Cape Town, Kremer’s model is resoundingly rejected – the coefficient on annual partners interacted with years sexually active is close to zero and tightly estimated for both men and women. However, the data also does not fit perfectly with to the simulated matching data; in particular, the coefficient on the marriage dummy is quite different.

Vuong’s (1989) test of model selection with non-nested hypotheses allows us to formalize which of the two models fits the data better. Creating a matched lifetime partners prediction and a Kremer prediction, I run a tobit of actual observed partners on each. That is, I examine the equation

$$Lifeprt_i = E_m [Lifeprt_i | X_i] + cons + \varepsilon_i$$

where  $E_m [Lifeprt]$  is predicted number of lifetime partners according to model  $m$  given data vector  $X_i$  and  $\varepsilon_i$  is normally distributed measurement error. In the case of Kremer’s model,  $X$  is a sufficient statistic for expected lifetime partners, however in the case of the matching model only the linear projection of simulated data onto available explanatory variables is utilized. As much of the variation of the matching model is thrown away, this is likely to bias results against accepting the matching model as a ”better fit.” Vuong’s test normalizes the ratio of likelihood ratio of these two models to have an asymptotic standard normal distribution, where a number significantly greater than zero states that the matching

hypothesis fits the data better, while a significant negative coefficient suggests that the taste-based model is more accurate. Vuong’s test provides strong support for the matching model over the taste-based one in both samples and for both genders.

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. Across cultures, the annual numbers of partners appear identical. In fact, even the difference in male versus female reporting appears to be quite constant between the two continents. Tables 13 and 14 explore other aspects of sexual behavior data, including age at first intercourse, and marital status for sexually active adolescents. The American data look nearly identical to the African data; indeed, Moreover, the two-type hypothesis which could also generate the steep age-death profile is also rejected – in both datasets, young adults report a continuous distribution of partners. Reported sexual behavior appears to remain remarkably similar across cultures.

## 9 Why Africa?

Intuitively, there is nothing about the marital search model which is unique to Africa. Indeed, two arguments have also been made to suggest that there is little about Africa which is idiosyncratic in terms of sexual behavior: first, other sexually transmitted diseases in different continents exhibit similar age-patterns of spread, and second, young people report strikingly similar behavior in America as in Cape Town. 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 reporting is often subject to hard to predict biases. However gonorrhea in the US provides a good case

study. Unlike HIV, gonorrhea is extremely infectious; with transmission probabilities very high for a single contact and approaching 1 for a month-long relationship. Gonorrhea 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 13 and 14 illustrate the predicted versus observed gonorrhea 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.

Behavior shared by Western and African cultures appears to drive the spread of sexually transmitted diseases. Why, then, has the US and Europe been spared a pandemic of the scale of that observed in Africa? Undoubtedly many factors play a role: for example, the transmission rate may be lower in the US due to a more circumcised male populace and lower (though still high) herpes prevalences. The quick and strong public health response to the first few AIDS cases may have encouraged Americans to more frequently use condoms with their short-term relationships<sup>14</sup>. However, the marital search model suggests an additional story. In order to achieve pandemic HIV, the marital search model demands a constant injection of a few percent who have been infected over many years, and small fluctuations in the size of this injection yield enormous differences in prevalence rates. The behavior of a tiny minority may well be the difference between Southern Africa and elsewhere. Indeed, Epstein (2002) has suggested that the difference between Southern and Eastern Africa could well be the culture of migratory labor in Southern Africa – migrants who undergo a search process both at home and at their worksite would fit the injection I describe.

---

<sup>14</sup>This has been offered as one of the explanations for the quick drop in gonorrhea prevalence observed in the 1980s, although a contemporaneous anti-gonorrheal campaign confounds analysis

## 10 Conclusions and Messages for Public Health Campaigns

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 is one of the few which passes 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 we shut down acute infection, prevalence rates fall dramatically. 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 using condoms for 3 months much more palatable than a lifetime of condom usage, and a tremendous amount of risk would be averted<sup>15</sup>. Indeed, this message contrasts strongly with the message adopted by many public health groups, who encourage condom usage 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.

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 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

---

<sup>15</sup>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 targeted risk seems likely to be one that individuals can understand.

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 more than 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 actual 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, not about why some countries reach very high prevalences and others do not. 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. This paper suggests that understanding differences in the epidemic's evolution in different contexts may require an understanding of the behaviors of small fractions of individuals, a topic which standard sampling methodologies are hard-pressed to deal with. Respondent-driven sampling method-

ologies (e.g. Heckathorn 1997) may be an asset in understanding the behavior of this pivotal minority.

## 11 Data Appendix

Three data sources are used in this article.

1) The Cape Area Panel Study<sup>16</sup> (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 and 2003 Kenya DHS surveys. DHS survey information is available at [www.measuredhs.com](http://www.measuredhs.com).

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](http://www.nrf.ac.za/sada)

3) 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 4928). More information is available at [www.cdc.gov/nchs/nsfg.htm](http://www.cdc.gov/nchs/nsfg.htm).

---

<sup>16</sup>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.

## 12 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

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

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

$$D(a, t) = \sum_{r=0}^a (S(r) - S(r+1)) (I^S(a-r, t-r) + I^M(a-r, t-r)) \quad (5)$$

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,  $\rho$ . In particular, let  $i^M(\mu, k, t)$  be the risk of infection for an individual who married  $\mu$  years earlier after  $k$  years of search in period  $t$ . then

$$i^M(\mu, k, t) = \psi^m(t - \mu) \rho (1 - \rho)^{\mu-1} \prod_{r=0}^k \left( 1 - \frac{\iota^s(t - \mu - r)}{1 - \psi^S(t - \mu - r)} \right) \quad (6)$$

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  $a - \mu$  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  $\tilde{\xi}^S(k)$  ( $\tilde{\xi}^M(k)$ ) 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

$$\frac{D(a, t)}{Popn} = \sum_{r=0}^a (S(r) - S(r+1)) * \left( \begin{array}{l} \gamma^S (1 - M(a)) \sum_{k=0}^a i^S(k, t) (X(a-k) - X(a-k-1)) + \\ \gamma^M M(a) \sum_{\mu=0}^a \sum_{k=0}^{a-\mu} i^M(\mu, k, t) (X(a-\mu-k) - X(a-\mu-k-1)) \end{array} \right) \quad (7)$$

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).  $\rho$ , 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  $\gamma^S$  and  $\gamma^M$ . 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,

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

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  $\gamma^S/\gamma^M$ , meaning that I can identify the death rate up to a constant. In year 15,  $D^S/D^M \cong 1.85^{17}$ . 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 following 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](http://www.assa.org.za). accessed February, 2005.
- 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](http://www.cdc.gov/std/stats/toc2003.htm). accessed August 27, 2005.

---

<sup>17</sup>The shape of the age-death curve is surprisingly insensitive to this constant, as the distribution of single risk looks fairly similar to the distribution of marriage risk suggested in the model.

- 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.
- 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.
- Epstein, H. (2002). The hidden cause of AIDS. *The New York Review of Books* 49(8).
- 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.
- 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 Acquired Immune Deficiency Syndromes and Human Retroviruses* 12(4), 427–28.
- 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>.
- Oster, E. (2005). Sexually transmitted infections, sexual behavior, and the HIV/AIDS epidemic. *Quarterly Journal of Economics* 120(2), 467–515.
- Pedraza, M.-A., J. Del Romero, F. Roldan, S. Garcia, M.-C. Ayerbe, A. R. Noriega, and J. Alcamí (1999). Heterosexual transmission 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.
- 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. Lallemand (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.
- 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.
- 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. Brahmbhatt, 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 (2001). South African Demographic and Health Survey 1998. [www.measuredhs.com](http://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](http://www.nrf.ac.za/sada).
- Te Velde, E. R. and P. L. Pearson (2002). The variability 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.
- 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 year 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

$$\begin{aligned}
&\theta \in \{0, 1, 2, \dots, 50\} \\
&F(\theta) = \text{Uniform}(0, 50) \\
&\rho_1^m = \rho_1^f = .2 \\
&\rho_2^m = \rho_2^f = .02 \\
&\Pr(\theta'|\theta) = \left\{ \begin{array}{l} .3|\theta' = \theta \\ .2|\theta' = \theta \pm 1 \\ .15|\theta' = \theta \pm 2 \end{array} \right\} \\
&T = 360 \\
&N = 100 \\
&\text{lifespan} = 240
\end{aligned}$$

Table 4: Parameter Inputs

	var	mean	sd	median	min	max
lifetime partners	13.306	9.955	11	1	67	
# partners, last 12 mos.	1.415	1.468	1	1	12	
total years single	8.646	6.303	9	1	19	

Table 5: Simulated Data Summary Statistics

		mean	std. dev	min	max
$\delta = .01$	all cohorts	0.0787	0.0175	0.0465	0.1116
	first 6 cohorts	0.0624	0.0158	0.0283	0.1
	last 6 cohorts	0.0925	0.0208	0.0529	0.13
$\delta = .02$	all cohorts	0.1445	0.02	0.1016	0.1929
	first 6 cohorts	0.1237	0.0203	0.0783	0.195
	last 6 cohorts	0.1706	0.0288	0.1114	0.2314
$\delta = .03$	all cohorts	0.2045	0.0204	0.1681	0.2494
	first 6 cohorts	0.1736	0.0255	0.1167	0.2317
	last 6 cohorts	0.2364	0.0293	0.1614	0.2857
$\delta = .04$	all cohorts	0.2544	0.0219	0.2155	0.3026
	first 6 cohorts	0.2233	0.0258	0.1683	0.2817
	last 6 cohorts	0.2866	0.0339	0.2186	0.37
$\delta = .05$	all cohorts	0.3037	0.0247	0.2445	0.3606
	first 6 cohorts	0.2636	0.0263	0.2117	0.32
	last 6 cohorts	0.34	0.0316	0.2586	0.3986
$\delta = .06$	all cohorts	0.3422	0.0172	0.3097	0.3832
	first 6 cohorts	0.3025	0.0265	0.2567	0.3583
	last 6 cohorts	0.3783	0.0267	0.3229	0.44

Table 6: Prevalence Results

Table presents simulated prevalences for each level of  $\delta$ , the inputted prevalence from all other sources.

		mean	std. dev	min	max
$\delta = .01$	all cohorts	0.0429	0.0065	0.031	0.06
	first 6 cohorts	0.0401	0.0102	0.0217	0.0667
	last 6 cohorts	0.0443	0.0105	0.0257	0.0771
$\delta = .02$	all cohorts	0.0826	0.0081	0.0603	0.099
	first 6 cohorts	0.0815	0.0134	0.055	0.11
	last 6 cohorts	0.084	0.0132	0.0529	0.1229
$\delta = .03$	all cohorts	0.1229	0.0099	0.1023	0.1426
	first 6 cohorts	0.1215	0.016	0.0933	0.16
	last 6 cohorts	0.1268	0.0157	0.0857	0.17
$\delta = .04$	all cohorts	0.1563	0.0099	0.1313	0.19
	first 6 cohorts	0.1535	0.0157	0.1233	0.1933
	last 6 cohorts	0.1587	0.0172	0.1157	0.2043
$\delta = .05$	all cohorts	0.189	0.0107	0.1632	0.2077
	first 6 cohorts	0.1869	0.0166	0.1417	0.2283
	last 6 cohorts	0.1875	0.0187	0.1529	0.24
$\delta = .06$	all cohorts	0.2222	0.0126	0.1955	0.2513
	first 6 cohorts	0.2216	0.0193	0.1783	0.2617
	last 6 cohorts	0.2203	0.0198	0.1786	0.27

Table 7: Prevalence Results: acute infection shut down

Reports simulated prevalence levels for each inputted  $\delta$  when acute infection is shut down

Variables are:	(1) Prev Pop Mn	(2) Prev Pop Mn	(3) Prev 40+ Mn	(4) Prev 40+ Mn	(5) HIV+ Individ	(6) HIV+ Individ
Age 1st sex	-0.033*** (5.288)		-0.008** (2.516)		-0.007*** (3.117)	
Age 1st Marriage	0.012** (2.173)		0.003 (0.913)		0.007*** (4.497)	
1st Marr-1st sex		0.018** (3.509)		0.004* (1.749)		0.007*** (4.632)
polygamous	-0.087 (1.365)	-0.013 (0.204)	-0.041 (1.592)	-0.032 (1.274)	0.040** (2.262)	0.039** (2.250)
spousal age gap	0.000 (0.040)	0.002 (0.925)	0.002 (1.374)	0.002* (1.810)	0.000 (0.103)	0.000 (0.156)
mnHIV					0.235*** (5.885)	0.234*** (5.906)
Observations	394	394	363	363	1847	1847
	OLS	OLS	OLS	OLS	Probit	Probit

Table 8: HIV and Duration of Singlehood

Columns 1-4 present ols regressions of Sampling Cluster prevalence on population mean characteristics, where in columns 3 and 4 only population means of individuals at least 40 years old are used as independent variables. Columns 5 and 6 present marginal effects on infection from an individual-level probit analysis, where mnHIV is the mean HIV prevalence among other individuals in the sampling cluster. t-statistics in Parentheses.

	OLS	OLS	
lifetime partners	CAPS - Males	Matching Sim	Kremer
Number of partners last year	0.3476*	0.8551***	1
	(1.75)	(5.31)	
num partners last year*(yrs sex active-1)	0.0312	0.1699***	1
	(0.66)	(4.15)	
years sexually active	-0.1196	0.4419***	0
	(0.85)	(3.85)	
constant	2.5435***	3.1283***	0
	(4.14)	(5.46)	
adj. R <sup>2</sup>	0.063	0.2109	
N	214	1000	
Vuong Stat	4.3448***		

Table 9: Comparison of CAPS data to Matching and Kremer models – Males  
Presents OLS regressions of lifetime number of partners on sexual behavior variables. Absolute values of t-statistics in parentheses.

	OLS	OLS	
lifetime partners	CAPS - Females	Matching Sim	Kremer
Number of partners last year	0.1172	0.8286***	1
	(0.80)	(6.11)	
num partners last year*(yrs sex active-1)	-0.0192	0.1599***	1
	(0.47)	(3.93)	
years sexually active	0.1023**	0.5144***	0
	(2.52)	(4.15)	
constant	1.4135***	2.7896***	0
	(7.07)	(5.28)	
adj. R <sup>2</sup>	0.0148	0.2324	
N	291	1000	
Vuong Stat	12.9038***		

Table 10: Comparison of CAPS data to Matching and Kremer models – Females  
Presents OLS regressions of lifetime number of partners on sexual behavior variables. Absolute values of t-statistics in parentheses.

	NSFG	tobit	tobit	Kremer
	lifetime partners	data	Matching Sim	
Number of partners last year		1.7256***	1.911***	1
		(17.48)	(16.10)	
num partners last year*(yrs sex active-1)		-0.0037	-0.0918***	1
		(0.46)	(9.77)	
years sexually active		0.1268***	0.043***	0
		(8.83)	(2.84)	
years single and sexually active		0.0850***	0.5200***	0
		(4.54)	(20.49)	
married		-1.8679***	1.5914***	0
		7.11	6.15	
married*years single and sexually active		0.1993***	0.3704***	0
		9.08	14.76	
constant		1.4697***	1.9318***	0
		7.64	7.84	
R <sup>2</sup>		0.1287	0.038	
Num Obs		2779	19000	
Vuong Stat		38.7739***		

Table 11: Comparison of NSFG data to Matching and Kremer models – Males  
Presents tobits of lifetime partners on annual partners and sexual behavior variables. For men, lifetime partners are right-censored at 7. Absolute values of t-statistics are in parentheses.

	tobit	tobit	
lifetime partners	NSFG Women	Matching Sim	Kremer
mon12prt	1.9156*** (9.16)	1.6845*** (25.97)	1
mon12prt*(ysa-1)	0.0779*** (4.31)	-0.0553*** (9.38)	1
ysa	0.0961*** (5.91)	0.052*** (3.44)	0
ysingle	0.2338*** (7.37)	0.5138*** (20.39)	0
married	-1.1841*** (3.03)	1.538*** (5.99)	0
married*ysingle	0.1831*** (5.19)	0.3505*** (14.13)	0
constant	0.1592 (0.39)	2.163*** (10.21)	0
R <sup>2</sup>	0.0342	0.042	
N	3094	19000	
Vuong Stat	5.58***		

Table 12: Comparison of NSFG data to Matching and Kremer models – Females  
Presents tobits of lifetime partners on annual partners and sexual behavior variables. For women, lifetime partners are right-censored at 50. Absolute values of t-statistics are in Parentheses.

NSFG Males						
stats	lifetime partners	Num partners last yr	Age first sex	ever had sex	married	
p10	1	1	13	0	0	
p25	2	1	14	0	0	
p50	4	1	16	1	0	
p75	7+	2	17	1	0	
p90	7+	4	18	1	0	
mean	4.0443	1.8106	15.8794	0.5474	0.0262	
sd	2.2865	1.6102	5.897	0.4974	0.1596	

CAPS Males						
stats		Num partners last yr	Age first sex	ever had sex	married	
p10		0	14	0	0	
p25		1	15	0	0	
p50		1	16	0	0	
p75		2	17	1	0	
p90		4	18	1	0	
mean		1.9914	15.7886	0.4577	0.0044	
sd		2.2518	1.8054	0.4983	0.0664	

Table 13: CAPS/NSFG Comparison - Males

Presents sample statistics for Males in the NSFG and CAPS data. p10, e.g., represents the tenth percentile in the data.

NSFG Females						
stats	lifetime partners	Num partners last yr	Age first sex	ever had sex	married	
p10	1	1	14	0	0	
p25	1	1	15	0	0	
p50	3	1	16	0	0	
p75	5	2	17	1	0	
p90	9	3	19	1	0	
mean	4.0667	1.4475	15.9475	0.4294	0.0935	
sd	4.75	1.219	2.0612	0.4952	0.2911	

CAPS Females						
stats		Num partners last yr	Age first sex	ever had sex	married	
p10		0	15	0	0	
p25		1	15	0	0	
p50		1	16	0	0	
p75		1	18	1	0	
p90		2	19	1	0	
mean		1.1321	16.5759	0.4612	0.0264	
sd		0.9386	1.7009	0.4985	0.1604	

Table 14: CAPS/NSFG Comparison - Females  
Presents sample statistics for Females in the NSFG and CAPS data. p10, e.g., represents the tenth percentile in the 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

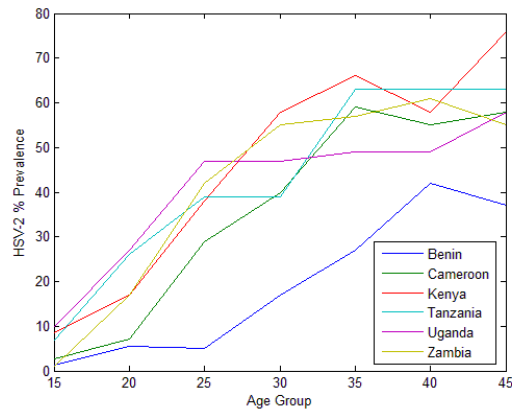


Figure 4: Male HSV-2 Prevalence by Age

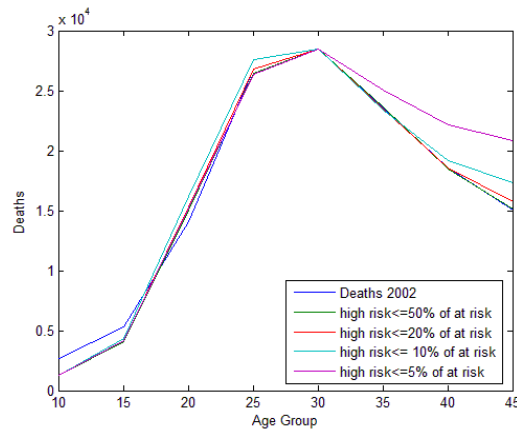


Figure 5: Deaths by Age: High Risk types  $\leq 20\%/yr$

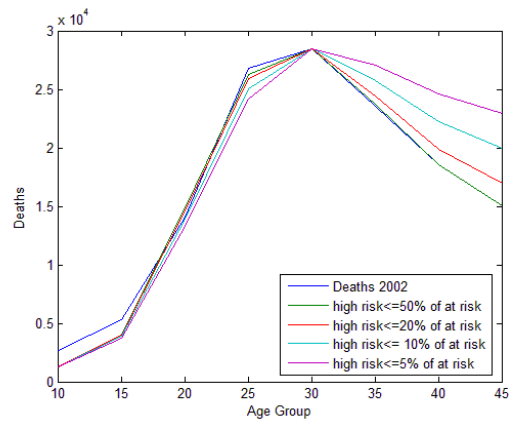


Figure 6: Deaths by Age: High Risk Types  $\leq 10\%/yr$




Figure 7: Deaths by Age: High Risk Types  $\leq 5\%/yr$




Figure 8: Best Fits at Various Epidemic Ages




Figure 9: 2002 Actual and predicted deaths by age for women in South Africa




Figure 10: 2002 Actual and Predicted Deaths by Age for Men in South Africa.




Figure 11: 2002 Actual, Matching, and ASSA women's deaths




Figure 12: 2002 Actual, Matching, and ASSA Men's Deaths




Figure 13: US Female Gonnorrhea by Age

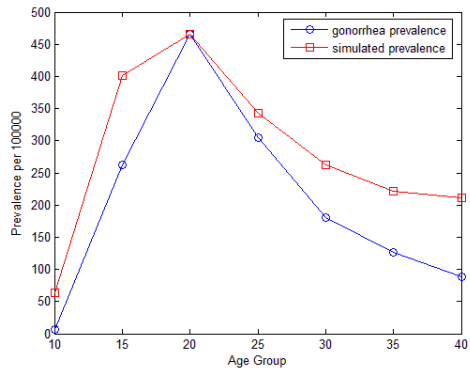


Figure 14: US Male Gonnorrhea by Age




Figure 15: Effect of Years Single on HIV