

An Equilibrium Model of the African HIV / AIDS Epidemic*

Jeremy Greenwood Philipp Kircher Cezar Santos
Michèle Tertilt

August 2018

Abstract

Twelve percent of the Malawian population is HIV infected. Eighteen percent of sexual encounters are casual. A condom is used a third of the time. To analyze the Malawian epidemic, a choice-theoretic general equilibrium search model is constructed. In the developed framework, people select between different sexual practices while knowing the inherent risk. The calibrated model is used to study several policy interventions. The analysis suggests that the efficacy of public policy depends upon the induced behavioral changes and equilibrium effects. The framework thus complements the insights provided by epidemiological studies and small-scale field experiments.

Keywords: ART, circumcision, condoms, disease transmission, epidemiological studies, HIV/AIDS, homo economicus, knowledge about HIV, Malawi, marriage, policy intervention, search, small field experiments, STDs, sex markets

JEL codes: I18, J12, O11, O55

**Affiliations:* Greenwood; University of Pennsylvania; Kircher, EUI and University of Edinburgh; Santos, FGV EPGE; Tertilt, University of Mannheim. Address correspondence to Tertilt at tertilt@uni-mannheim.de. We thank Ufuk Ackigit, Marco Basetto, Paco Buera, Isabel Correia, Mariacristina de Nardi, Pascaline Dupas, Theresa Fedor, Pedro Ferreira, Erik Hurst, Olga Itenberg, Greg Kaplan, Katja Kaufmann, Corinne Low, Robert Shimer, Nancy Stokey, Alessandra Voena, Mark Wright, and seminar audiences at numerous institutions for helpful comments. Financial support from NSF grant SES-0748889, ERC grant SH1-313719, and the Alfred P. Sloan Foundation is gratefully acknowledged. Tertilt thanks the Becker-Friedman Institute at the University of Chicago and the ifo Institute in Munich for their hospitality when parts of this project were completed. Last but not least, Hanno Foerster, Vera Molitor, Luciene Pereira, Xiaodi Wang, Xue Zhang, and Guilherme Zimmermann provided excellent research assistance.

1 Introduction

HIV/AIDS is a major cause of death, currently killing about 1 million people worldwide each year. The number of new infections is about 2 million, suggesting an even more severe problem in the future. The most affected continent is Africa, which hosts about two thirds of all HIV/AIDS infected people. Within Africa most transmissions occur through heterosexual sex. Furthermore, the majority of the HIV-positive population is female, compared to less than one third in most developed countries.

A natural question is what can and should be done to prevent the disease. Several policies have been at the center of discussion—antiretroviral therapy (ART) has been developed to keep infected people alive longer but also to reduce their infectiousness to others. The World Health Organization advocates large-scale male circumcision. Treating other sexually transmitted diseases (STDs) has been advocated as well, and research is under way to find a suitable vaccine. Others advocate the “ABC approach”—Abstinence, Being faithful, and Condom use.

The overall effect of these approaches depends not only on their medical efficacy, but also on the conduct of the population: A treatment that has been medically shown to reduce HIV transmission might be less effective if individuals start engaging in more risky sexual practices. This behavioral response gets amplified in equilibrium since less infections of others reduces the negative consequences of risky sexual practices further. The idea that an endogenous response by the population could reduce policy effectiveness has long been recognized in the theoretical disease transmission literature (see Philipson and Posner (1993)). In the most extreme case behavioral adjustment can be so large that a policy backfires, leading to more HIV (as in Kremer (1996)). Yet, just how much endogenous reactions might quantitatively reduce the effectiveness of interventions or even negate them has not been assessed to date. In fact, existing theoretical models are not designed for a quantitative assessment. Quantitatively accounting for this behavioral channel seems prudent so as to not overstate the likely effectiveness of policies, and this paper aims at providing a step to fill this gap.

The main building block is a novel computational choice-theoretic equilibrium

model of sexual behavior. The model features several margins of risky sexual behavior with two goals explicitly in mind: 1) to capture aspects of choice that are particularly relevant to HIV transmission, and 2) to use them to match against existing data—an aspect that was not possible in previous theoretical work.¹

In the tradition of *homo economicus*, each individual in the model rationally chooses their sexual behavior to maximize expected discounted life-time utility. Specifically, people choose how hard to search for a sexual partner. They can search in three different “markets” or “meeting places” for three different types of relationships — long term, casual sex using condoms, and casual unprotected sex. This market structure conveniently eliminates any problem with differences in interests between partners: they have the same desires when they meet in the same market. Finding a partner generates utility from sexual behavior. Marriage has the additional benefit of continued interaction without the need to search again. Searching in a market has both a convex effort cost and the cost of possibly contracting HIV. In line with the data, transmission rates differ according to gender, condom usage, ART availability, and male circumcision. Now, people know whether or not they are infected with HIV. But, crucially, when agreeing to sex, individuals do not know whether their partner is infected, treated or circumcised. When entering a meeting place for a relationship, they rationally recognize that some of these markets will be riskier than others.²

Men and women feature separately in the model. To capture heterogeneity in society, the degree of patience is allowed to vary across individuals, which induces different people to weigh the risk of sexual behavior differently. As people age, they become on average more patient. Men may be circumcised or not. Peo-

¹The economics literature on HIV transmission is small. The discussion in Philipson and Posner (1993) is mostly verbal. The seminal theory work in Kremer (1996) abstracts from important margins such as gender asymmetry, condom use, and timing considerations. Similar approaches feature in recent theoretical work on general disease transmission and decision-making, without focus on the specifics of HIV (e.g., Klein et al. (2007), Fenichel et al. (2011), Toxvaerd (2010)). An exception is Magruder (2011) who builds a Jovanovic (1979) style matching model of marital partner search and HIV, but abstracts from condom use and equilibrium effects; in his analysis choices are not affected by the level of the disease in society. While the focus here is on the disease itself, several people have analyzed the importance of HIV/AIDS for development (Young 2005; Santaaulalia-Llopis 2008).

²This echoes a theoretical point already made by Kremer (1996) that endogenous sorting can render sex in marriage safer than casual sex.

ple are endogenously heterogeneous in whether they are healthy or HIV-infected (with or without symptoms). In the presence of ART, infected people will further differ by whether or not they are being treated.

A stationary equilibrium is solved for. In equilibrium, each market is characterized by its endogenous riskiness and by a transfer that one partner makes to the other. These transfers clear the market and depend on the desires of men relative to women. There is also an adverse selection problem: Individuals with a proclivity toward risky sexual behavior are inclined to enter the market for casual unprotected sex. As a consequence, this market tends to have a high rate of HIV. The private information problem regarding a partner's health status dissuades healthy individuals without such a proclivity toward risky sexual behavior from entering this market. This exacerbates the riskiness of the market for casual unprotected sex. The presence of multiple markets, private information, and randomness in meeting a partner and in contracting HIV renders the model too complicated for analytical results. Instead a parameterized version of the model is used for quantitative analysis. The numerical benchmark is calibrated to match key moments of HIV and sexual behavior in Malawi. The calibration matches most targeted moments well, even though the model is sparse on gender differences. It also captures patterns in non-targeted life-cycle moments.

Policy insights and further support for the model come from the study of two policies aimed at reducing transmission risk: male circumcision and the treatment of other STDs. Medical research attributes a reduction in transmission risk to both interventions. The calibrated model predicts that both are effective in curbing HIV prevalence.³ The beneficial effects of reduced HIV transmission powerfully accumulate in equilibrium and dominate the (nevertheless non-trivial) behavioral adjustments going in the opposite direction.

These interventions also allow for further assessment of the model and show its ability to reconcile some outstanding puzzles in the literature. First, for circumcision, findings from small-scale field experiments are used to calibrate the

³Throughout the paper "HIV rate" is used as shorthand for "HIV prevalence rate." Note that prevalence is a stock variable, while one might also be interested in the flow; i.e. new infections. Whenever incidence (measured as new infections per 1,000 healthy people) is meant instead, this is explicitly specified.

reduction in transmission risk for an individual. These studies have observed changes in the sexual behavior of treated individuals, as is also found here. If one just applied the reduced transmission risk and behavioral change at the individual level, without accounting for the spillover effects on their partners, the HIV rate would only be moderately affected by circumcision. Yet in cross-country data—which serve as the next-best assessment in lieu of large-scale circumcision experiments—one sees large changes in the HIV rate with respect to the number of circumcised males. At the country level, equilibrium forces kick in as not only individuals but also their partners are affected. The model replicates the (non-targeted) cross-country evidence on circumcision, yielding credibility to the elasticities that the model produces. For STDs one of the puzzling aspect is that medical research indicates that effects should be positive, but most field experiments do not find effects. This absence of effects is reconciled in a partial equilibrium version of the model that mimics a field experiment: the lack of compounding and the additional risk-taking offset most effects.

ART was introduced in Malawi in 2005 and is also examined here. By 2014 about 50% of HIV-infected individuals were on ART. Given the fall in the HIV rate at the time, the Malawian government argued that the ART campaign has been a success. The results indicate that a causal connection is unlikely. ART is a complex policy. By treating the sick, infected people live longer and have more time to infect others. Moreover, being sick is no longer perceived as such a bad event and hence healthy people engage in more risky behavior. Even though the treated are less infectious to others, the former effects may dominate and HIV may actually rise. The quantitative results show that whether ART reduces or increases the HIV rate is a function of the fraction of people treated. When less than 50% of the infected are treated, the two effects roughly cancel out, and the HIV rate is largely unaffected. Only when more than half of the infected are treated, the reduced infectiousness becomes powerful enough to decrease the overall HIV rate.

This paper focuses mainly on “hard” interventions that directly target disease transmission. An intermediate intervention is one component of the ABC approach. Suppose one could increase the utility from condom usage, either technologically by improving their quality and comfort or through soft interventions

that reduce their stigma.⁴ The simulations suggest that better condoms have a potential to backfire. People would use them, but they would in parallel increase their sexual activity, which could lead to an increase in the HIV rate.⁵

Finally, the paper analyzes the importance of information for the disease. An extension is modeled in which a fraction of the population is not aware that abstinence and condom use reduces the odds of getting infected. The extension is used to investigate to what extent better information contributed to the HIV decline over time. Reducing the fraction of uninformed people in line with the data, the model predicts a decline in HIV incidence between 1996 and 2004 comparable to the one observed in the data. The reason is that better informed people engage in less risky behavior. This is also in line with the data. Thus, better information probably played a large role in the Malawian success story.

This work ties in with two strands of literature. The epidemiological literature on disease transmission in general, and HIV in particular, is large, but tends not to model decision-making and takes sexual behavior as exogenous. However, several studies suggest that people react to a higher presence of HIV/AIDS by adjusting aspects of their sexual behavior [Wellings et al. (1994)]. One contribution of the present paper is providing a way of accounting for the adverse incentives for individuals to increase their risky practices. It thus provides a more cautious view on the effectiveness of interventions. Viewed through the lens of the model, omitting behavioral adjustment would lead to limited bias in the case of circumcision and would substantially overstate the effectiveness of treating STDs and of ART. The workhorse epidemiological model of disease transmission is the susceptible-infected (SI) model with random mixing; see, e.g., Anderson and May (1992).⁶ In such a model people are either infected or susceptible. If a person

⁴Acceptability of condoms is a major issue, as reviewed in the case of South Africa in Bekinska, Smit, and Mantell (2012). They cite studies that attribute low acceptability to a number of hard technical problems such as “comfort and breakage” and “unpleasant smell”, but also to soft problems of negative attitudes due to stigmatization because of the link to STDs.

⁵The discussion of other purely soft interventions aimed not at the underlying transmission channel but at changing attitudes (e.g. towards marriage) are relegated to the companion note; see Greenwood et al. (2017b).

⁶Many papers simulate SI models numerically to forecast the disease incidence or to study the effectiveness of prevention policies. For example, Low-Beer and Stoneburner (1997) use such a model to forecast HIV incidence in East Africa; Johnson (2008) studies the role of other sexually transmitted diseases for the epidemic; Clark and Eaton (2008) simulate the effectiveness of male

is infected, they transmit the disease to susceptible (non-infected) people until the individual leaves the sexually active population. In models of HIV/AIDS there is no stage of recovery.⁷ In these models people encounter other individuals in the population randomly. An interesting contribution is Kremer and Morcom (1998) who introduce selective mixing into a SI model. This idea also features prominently here, but is applied to the type of sexual behavior that a person seeks. The special case where people exclusively meet others who are seeking the same type of relationship allows for selective mixing and avoids modeling any conflict of interest in relationship formation.

The second large literature to which the current work connects is the recent empirical literature that studies HIV/AIDS prevention policies using mostly field experiments (see Padian et al. (2010) for a survey). While some are large in scale, the vast majority treat only a small fraction of the population. The absence of large randomized studies of male circumcision has been strongly voiced (see e.g., Williams et al. (2006) and De Walque et al. (2012)).⁸ Small studies establish the effect of a treatment on an individual, without reshaping the rest of the market. These insights are very valuable and are used in the present study to calibrate some parameters. Such studies are not designed to assess the effects when everyone in a population is treated, as they omit two effects. First, on the positive side, treatment affects not only individuals, but also their partners, and their partners' partners, and so on. Second, on the negative side, people might become even more reckless if their partners are deemed safer. Through the lens of the model, individual effects are substantially smaller than equilibrium effects for several policies. This positions the current paper as a way to extend insights from small-scale field experiments to full equilibrium while being prudent with respect to behavioral responses. The importance of accounting for behavioral responses is indeed a lesson from existing field experiments: Auld (2006) an-

circumcision; and Bracher, Santow, and Watkins (2004) study the importance of condoms. By design, none of these studies takes behavioral changes, in response to the reduced transmission risks, into account.

⁷Recovery, and either resistance against further infection, or the possibility of reinfection, are explicitly modeled for other infectious diseases. See Hethcote (2000) for a comprehensive overview of the mathematical modeling of infectious disease.

⁸For examples of large-scale investigations, several broad studies find that male circumcision lowers the infection risk for males [see, e.g., Auvert et al. (2005) and Gray et al. (2007)], but they are not randomized.

alyzes, using data from San Francisco in the 1980s, how risky sexual behavior changes with local infection rates and finds large elasticities. Lakdawalla, Sood, and Goldman (2006) study the effects of ART and find an increase in risky sexual behavior in the United States. Dupas (2011) shows in a randomized field experiment in Kenya that teenage girls who are given information about the HIV status of different groups of men respond by shifting sexual behavior to the lower risk groups.⁹

A caveat is in order before proceeding. Research using computational general equilibrium models to assess the implications that interventions might have on the spread of HIV/AIDS (or other diseases) is in its infancy. Overall, this research program aims to develop tools to aid researchers and practitioners and highlights areas where further and more in-depth research should be conducted.

The paper is organized as follows. Section 2 provides information on sexual behavior and HIV/AIDS in Malawi. Section 3 sets up the model and defines the equilibrium. Section 4 describes the benchmark calibration. Section 5 presents the results of the policy experiments. In the early stages of the HIV/AIDS epidemic the population may not have been fully aware of the transmission process. Section 6 discusses the role that better information could have had on mitigating HIV. Section 7 offers some conclusions.

2 Families, Sexual Behavior, and HIV/AIDS in Malawi

The Republic of Malawi serves as a focal country to which the analysis is applied. Therefore, this section briefly describes some information on the HIV/AIDS epidemic in Malawi, together with details about sexual behavior and family life. This background will be useful in guiding the modeling choices.

The Republic of Malawi is a country in southeast Africa with a population of 14

⁹Other insights that have emerged in the context of HIV are: testing is not a very cost-effective prevention policy in a randomized field experiment in rural Malawi (Thornton (2008)); HIV/AIDS information campaigns have a larger impact on more educated people (De Walque (2007) , but see also the analysis in Iorio and Santaaulalia-Llopis (2011)); non-financial incentives are a good motivator to induce hairdressers to also sell condoms as a prevention policy (Ashraf, Bandiera, and Jack (2014)).

million people. Malawi suffers greatly from the HIV/AIDS epidemic. In 2004, twelve percent of the adult population was infected.¹⁰ This is well above the average within Sub-Saharan Africa (S.S.A.), which has an adult prevalence rate of about 7.2%—see Canning (2006). It is also well below the HIV rate of the most affected countries, such as Botswana with an adult prevalence rate of 37%. Like in many other countries, HIV in Malawi displays a hump-shaped profile over time. The HIV prevalence rate peaked in 1999 and has been continuously falling since. The incidence rate peaked earlier, in 1994 with 21.8 infections per 1,000 healthy population.

The principal mode of HIV transmission in Malawi is through heterosexual sex. Mother-to-child transmissions account for about 10% of all new HIV infections. This fact is ignored here. Most people born with HIV die before they reach sexual maturity and therefore do not add to the propagation of HIV. Like the rest of Sub-Saharan Africa, more than half of the HIV-infected population in Malawi are women. By contrast two thirds of the infected population are men in the Western world—see *World Development Indicators* (2009). The HIV rate among adult women is currently about 13%, compared to 10% among men, suggesting important gender differences.

A rational behavior model of HIV only makes sense if people understand what HIV is, are aware of how it gets transmitted, and know how to avoid it. This seems largely to be the case in Malawi since the 2000s. Almost 100% of surveyed Malawians had heard of HIV or AIDS. About 57% of women and 75% of men correctly identified the use of condoms as a means to protect against HIV infection. Finally, an overwhelming majority of adults in Malawi—74% for women and 86% for men—know where to get condoms. Further, Delavande and Kohler (2009) document that people in Malawi are relatively good in assessing their own probability of being infected with HIV. Prior to 2000 people were less informed about the HIV transmission mechanism. Due to ignorance they might not have switched toward safer sexual practices. The analysis here examines whether the diffusion of information about HIV could have contributed to the observed decline in the Malawian HIV rate as people engaged in less risky sexual behavior.

¹⁰Unless noted otherwise, information on HIV prevalence and patterns of sexual behavior are from the 2004 Demographic and Health Survey's (DHS) Final Report for Malawi.

Treating HIV with antiretroviral drugs was introduced not until the mid 2000s in Malawi. Treatment increased gradually from 3% of the infected in 2005 to 50% in 2014. The importance of treatment will be examined in detail in this paper. About 20 percent of the male population were circumcised in Malawi in 2004. Circumcision seems largely related to religion and ethnicity rather than health concerns, which explains its relative stability over time (in 2010, 22% of men were circumcised).¹¹ There has been a slight uptake more recently, likely related to official HIV prevention interventions, so that by 2015, 28% of men were circumcised.

Sexual behavior conducive to the spread of the disease is relatively common in Malawi. Condoms are used by less than half of all respondents in their last sexual act. It is also considered normal for unmarried people to change partners often, see Undie, Crichton, and Zulu (2007). Furthermore, divorce is relatively common. Reniers (2003) reports that 45% of marriages end in divorce within 20 years. Several other forms of risky behavior will be abstracted from in the paper. For example, the model advanced here does not allow concurrent multiple relationships, such as extramarital affairs or polygyny. The model abstracts from concurrent relationships to keep it tractable. Future work should include such relationships.

The high prevalence of risky behavior does not necessarily imply that people are uninformed or irrational: it is more likely due to the trade-off between increased safety versus less pleasure. For example sex with a condom is often compared to eating candy “while it’s in the wrapper” or a banana with its peel (Undie, Crichton, and Zulu 2007). Condom use within marriage is essentially non-existent in Malawi [Chimbiri (2007)]. One reason is that marital sex is often aimed at reproduction. Furthermore, using a condom in marriage may be interpreted as a signal of infidelity [Bracher, Santow, and Watkins (2004)]. Note that while using a condom lowers transmission risk substantially, it does not decrease the risk to zero. Bracher, Santow, and Watkins (2004) cite a study that finds that for new condoms, the average breakage rate is 4%; this rate jumps to 19% for condoms that are 7 years old.

Poulin (2007) documents that money and gift transfers in sexual partnerships

¹¹For example, 93% of Muslims are circumcised but only 8.6% of Catholics. Similarly, 82% of the Yao ethnicity are circumcised but only 2% of the Tumbuka. See DHS 2004.

are part of courting practices in Malawi (see also Swidler and Watkins (2007)). In addition to an expression of love and commitment, she argues that these transfers are a way of acquiring sex for men and about meeting their financial needs for women. A gift might be in the form of sugar or soap, but also in cash. Transfers are not made directly before or after sex (as with prostitution), however; rather gift giving is an integral part of a relationship. The model developed here allows for such transfers between men and women in sexual relationships.

3 Economic Environment

Imagine a world populated by males and females. Males and females desire relationships with the opposite sex. There are two types of relationships, viz. short-term and long-term ones. Within a relationship individuals engage in sex. Sex is risky because of the presence of the HIV/AIDS virus in society. There are two types of sex, protected and unprotected. Protected sex offers a better defense against the transmission of HIV/AIDS across partners. It provides less enjoyment, though. Individuals interested in a short-term relationship must decide what type of sex they desire. Put simply, they must weigh the extra momentary utility associated with unprotected sex against the increased odds of getting infected with the HIV/AIDS virus. As motivated in Section 2, sex is always unprotected in long-term relationships. Further, suppose that a person can only engage in one relationship at a time.

Denote the utility from unprotected sex by u and the utility from protected sex by p , with $u > p > 0$. The utility flow in a long-term relationship is $u+l$, where l may be negative. A positive l can be interpreted as a taste for long-term attachment, while as a negative l can be construed as taste for variety in partners.¹² Individuals also realize utility from the consumption of goods. Let this utility be given by $\ln(w)$, where w is consumption (“wealth”). Each period a person receives income in the amount y . There is no borrowing or saving in the economy.¹³ An individual discounts the future with a stochastic factor that takes two values, viz. \tilde{t} and

¹²In section 5.3.1 an extension is presented in which married people derive an additional utility flow from having children.

¹³Magalhães and Santaaulàlia-Llopis (2018) provide evidence for the lack of savings in Malawi.

$\tilde{\beta}$ with $\tilde{\iota} < \tilde{\beta}$. There is a distribution across individuals for the $\tilde{\iota}$ s.¹⁴ Connected with each $\tilde{\iota}$ is a $\tilde{\beta}$. A person starts off life with the low rate $\tilde{\iota}$ (dubbed “young”). This low factor reflects the impatience of youth, which results in a predilection to engage in risky behavior.¹⁵ Then, in every period, a person may switch permanently to the high factor with probability η (dubbed “old”). Additionally, there is a probability δ that an individual dies from natural causes in a period. Thus, the effective discount factors are given by $\iota \equiv \tilde{\iota}(1 - \delta)$ and $\beta \equiv \tilde{\beta}(1 - \delta)$. So, the pair (β, ι) represents the sample space from which an individual draws his effective discount factor. This sample space differs across people, as noted above. Finally, a male individual may be circumcised (denoted by $c = 1$) or not ($c = 0$). A circumcised male is less likely to contract the HIV virus from his sexual partner. The values of c , β , and ι differ across individuals of a given gender. The set of fixed characteristics for a person is denoted by $x = (c, \beta, \iota)$, called a person’s type.

People can search for partners in different markets.¹⁶ Depending how much they value a particular relationship, people choose their search effort by picking the odds of finding a partner which comes at a cost. At the beginning of each period an unattached individual may search for a long-term partner. The odds of finding a partner on the long-term market are denoted by π_l . The individual picks these odds at an increasing cost in terms of lost utility. These search costs are given by $C_l(\pi_l) = \omega_l[\pi_l/(1/2 - \pi_l)]^{\kappa_l+1}$, where $\kappa_l \geq 0$ and $\omega_l > 0$. Observe that $C_l(0) = 0$ and $C_l(1/2) = \infty$. A long-term relationship may break up (at the end of) each period with exit probability ϵ . If the person is unsuccessful at finding a long-term mate, then they enter the short-term market, where they can still engage in sexual behavior for this period. Note that an individual who does not want a long-term relationship can set $\pi_l = 0$. If the person wants casual sex, they can choose the search effort for protected and unprotected sex simultaneously. Let π_p and π_u

¹⁴Heterogeneity in discount factors leads to people differing in their risk behavior. Some risky people are needed to generate any HIV in equilibrium. The importance of risky groups (e.g., truck drivers and prostitutes) is discussed in the literature (e.g., Kremer (1996), Nzyuko et al. (1997) and Oster (2012)).

¹⁵The impatience of youth and the resulting higher risk-taking behavior has been documented in many contexts. For example Rolison et al. (2013) find that risk-taking attitudes reduce smoothly with age, especially in the health domain.

¹⁶The idea that people can rationally target their search behavior to particular markets is present in many recent theoretical models (e.g. Jacquet and Tan (2007), Eeckhout and Kircher (2010), Gautier, Svarer, and Teulings (2010)).

represent the odds of finding a partner in the protected and unprotected markets for short-term relationships, which are choice variables. If she prefers one type of sex over the other, she can choose the effort (and hence the odds) accordingly. The cost of searching in each market is given by $C_s(\pi_p)$ and $C_s(\pi_u)$, which have the same functional form as $C_l(\pi_l)$, but where the parameters κ_s and ω_s are allowed to differ from the long-term market. The total cost of searching for a short-term partner will then be $C_s(\pi_p) + C_s(\pi_u)$. An individual cannot simultaneously draw a partner on both markets. Since $C_s(1/2) = \infty$, the odds are such that $\pi_p + \pi_u < 1$, and an individual will be abstinent with probability $\pi_a \equiv 1 - \pi_p - \pi_u$. Also, observe that individuals can choose abstinence by picking $\pi_p = \pi_u = 0$.

The probabilistic nature of finding a partner in the three markets is meant to capture the randomness in everyday life of meeting someone. Finding a partner is costly. The more you invest in it, the more likely you are to succeed. While π_l , π_p , and π_u represent the probabilities of an individual finding a partner in the three markets, they also represent the fractions of people searching in each markets that will find a partner, given the large nature of the economy. Last, this randomness in the meeting structure helps to smooth out the equilibrium under study because people who are currently of the same type will evolve differently depending on the relationship they enter into.

Given the pervasive evidence on gift giving in the context of sexual relationships (see Section 2), transfers are exchanged for sex. Associated with each market is a transfer payment, z , that is made between the two partners. For the person making the transfer, z will be positive, while it will be negative for the individual receiving it. Think about the people receiving the transfers as supplying relationships on the market, and those paying transfers as demanding them. The magnitude of this transfer is determined in equilibrium.¹⁷ It will depend upon the demand and supply for a given type of relationship by each gender. This hinges on the utility that each gender realizes from a partnership in the various markets and the riskiness of participating in them.

People know their own health status ϕ .¹⁸ A healthy individual has $\phi = 1$. An

¹⁷Alternatively, market clearing could be achieved through different meeting probabilities. This should lead to qualitatively similar results.

¹⁸Greenwood et al. (2013) carefully explored the evolution of beliefs about the own likelihood

individual with HIV infection and no antiretroviral therapy (ART) treatment has $\phi = 0$. An infected individual who receives treatment has health status $\phi = t$. So, $\phi \in \{0, 1, t\}$. All individuals are born healthy. If an individual has sex with a partner of health status $\hat{\phi}$, then the virus will get transmitted with probability $1 - \gamma(\hat{\phi})$, which is trivially zero if the other individual is healthy.¹⁹ Further, $1 - \gamma(t) < 1 - \gamma(0)$, so that a treated individual is less likely to infect others. Similarly, an individual of type ϕ transmits the infection to a healthy partner with probability $1 - \hat{\gamma}(\phi)$. The transmission probabilities can differ by gender. Circumcised men are also less likely to contract the virus compared to non-circumcised individuals. Denote this reduction in transmission probabilities that circumcision provides by $\chi(c)$, with $\chi(1) = \chi < 1$ and $\chi(0) = 1$. For women, trivially, $\chi_f(c) = 1$. These transmission probabilities are also lower for protected sex than for unprotected sex. People only know their own health status. While they cannot discern the health status of other individuals, they hold correct expectations, $R_r(\hat{\phi})$, about the fraction of potential partners of each health status in each market, $r = l, p, u$. (As will be discussed later, for women $R_i(\hat{\phi})$ will also depend on whether or not their partner is circumcised.)

If treatment is available, a currently infected untreated individual ($\phi = 0$) will enter next period under ART with probability q . Once in treatment, the individual remains in treatment forever. So the probability of obtaining treatment next period, conditional on current health status, ϕ , can be summarized by the function $Q(\phi)$, where $Q : \{0, 1, t\} \rightarrow \{q, 0, 1\}$; i.e., a currently infected untreated person obtains treatment next period with probability $Q(0) = q$, a healthy person will not receive treatment corresponding to $Q(1) = 0$, and a treated person will still receive ART implying $Q(t) = 1$.

The health status indicator refers to individuals in the early stages of the disease, where it is assumed that their health status is not visible to other people and neither income nor sexual activity is restricted. Individuals transit to the final stages of AIDS (where they display symptoms) with probability α_ϕ , where

of being HIV positive in a world where individuals do not immediately observe their own health status. People update their beliefs about being healthy in a Bayesian fashion, conditional on their past sexual history (and also on what they observe about their partners in long-term relationships). An editor is thanked for suggesting the current simplification.

¹⁹The symbol $\hat{\cdot}$ denotes the characteristics of an individual's partner.

$\alpha_1 = 0 < \alpha_t < \alpha_0$ since healthy individuals do not enter the final stage of the disease and treatment prolongs a healthy existence. Assume that a person stricken with final stage HIV/AIDS symptoms engages in no further relationships. Let the remaining lifetime utility for a person with the symptoms of HIV/AIDS be represented by A .²⁰ The probability that a person displaying symptoms dies is δ_2 . Since a person with HIV/AIDS symptoms engages in no further sexual activity, δ_2 does not appear in the value functions. It is still relevant for computing the average HIV/AIDS rate in society.

Note that in the framework there is attrition in the population each period both due to natural death and to HIV/AIDS. This loss is replenished by an inflow each period of newly born males and females. Assume that $\mu(x)$ type- x individuals are born at the beginning of each period. Recall that x denotes the set of permanent characteristics for an individual, namely c, β , and ι . People also differ by gender. Gender will be suppressed unless it is specifically needed and then it will be represented by the subscript g (for $g = f, m$) attached to a function or variable.

Before proceeding on to the formal analysis some notation will be defined. An individual will be indexed by his health status, $\phi \in \{0, 1, t\}$, his current discount factor, d , drawn from his fixed sample space $\{\iota, \beta\}$, and his exogenous type, $x = (c, \beta, \iota)$. Let $\tilde{V}_r^d(\phi, x)$ denote the expected lifetime utility for a person with health status ϕ , realized discount factor, d , and exogenous type x , who just found a partner for a relationship of type $r = a, l, p, u$ (abstinent, long-term, short-term protected and short-term unprotected). Similarly, $V_r^d(\phi, x)$ will represent the expected lifetime utility for a person who is currently searching for a partner in a type- r relationship (for $r = l, s$ where s denotes short term), but has not found one yet. The timing of events is summarized in Figure 8 in Appendix A. Attention will now be directed toward the determination of the functions $\tilde{V}_r^d(\phi, x)$ and $V_r^d(\phi, x)$. The focus is on studying a stationary equilibrium.

²⁰Note that A cannot be too large to ensure that people prefer being healthy over being sick.

3.1 Short-term Relationships

3.1.1 Abstinence

The case of abstinence is the easiest to analyze. Recall that there are young and old individuals that differ in their discount factor. Start with a type- x old person (i.e., with discount factor β), with health status ϕ , who has failed to match on the short-term sex markets. Thus, the person will be abstinent in the current period. Note that the individual's discount factor will remain high forever. The value function for this person is given by

$$\tilde{V}_a^\beta(\phi, x) = \ln(y) + \alpha_\phi \beta A + (1 - \alpha_\phi) \beta \{Q(\phi) V_t^\beta(t, x) + [1 - Q(\phi)] V_t^\beta(\phi, x)\}, \quad (1)$$

for $\phi = 0, 1, t$. The first term on the right-hand side covers the flow utility from current consumption. The continuation value depends on whether the person enters the final stages of the disease, which happens at rate α_ϕ and yields a discounted continuation value of βA . With complementary probability, $1 - \alpha_\phi$, the individual's continuation value is the discounted expected value of either continuing in their current state or of perhaps obtaining treatment. Healthy individuals have a zero probability of entering the final stages of the disease or of obtaining treatment, so their value function reduces simply to $\tilde{V}_a^\beta(1, x) = \ln(y) + \beta V_t^\beta(1, x)$, where the discount factor already incorporates other sources of death apart from HIV/AIDS.

Next, consider the case of an abstinent young person (with discount factor ι). The discount factor may switch next period to the high value, β , with probability η , or remain at the low one, ι , with probability $1 - \eta$. Therefore, the value functions for young individuals retain the same structure as (1) with two modifications: first, the discount factor β needs to be replaced by ι and, second, the value function $V_t^\beta(\phi, x)$ on the right-hand side has to be replaced by the expected value $\eta V_t^\beta(\phi, x) + (1 - \eta) V_t^\iota(\phi, x)$. This difference between the value functions for young and old individuals holds throughout the analysis. Therefore, the focus is placed on defining the value functions for old individuals, with the adjustments for young individuals summarized in Appendix D.1.

3.1.2 Sexual Relationships

Now consider short-term relationships. Here individuals have to take into account both the transfers in each market, as well as the distributions over health status of partners they might encounter. Additionally, they also experience the utility from sexual activity. For now, assume a person is already matched. The search behavior for an unmatched individual will be described at the end of this section.

Again, focus on an old individual with a high discount factor, β . If $s = p$, then the person will use a condom and enjoy utility p from sex. If $s = u$, the individual will enjoy u from unprotected sex. Define the indicator function $I(s)$ to return a value of 1, when $s = p$, and a value of 0, otherwise. Thus, the flow joy from a short-term sexual relationship can be written as $pI(s) + u[1 - I(s)]$. Apart from this addition and the incorporation of the transfer z_s in the short-term market, the value function of infected individuals with $\phi \in \{0, t\}$ looks similar to the one for abstinence (1):

$$\begin{aligned} \tilde{V}_s^\beta(\phi, x) = & \ln(y - z_s) + pI(s) + u[1 - I(s)] + \alpha_\phi\beta A \\ & + (1 - \alpha_\phi)\beta\{Q(\phi)V_t^\beta(t, x) + [1 - Q(\phi)]V_t^\beta(0, x)\}, \end{aligned} \quad (2)$$

for $s = p, u$. It comprises the utility from current consumption and sexual activity, and the continuation value of either entering the last stages of the disease or from continuing as an infected or a treated individual.

A healthy individual ($\phi = 1$) never enters the last stages without having been in the infected/treated status for a period, but they have to weigh the costs of various sexual activities in terms of their chances of getting infected. First, the transmission risk of catching HIV/AIDS from an infected person differs across markets. Specifically, the transmission risk in the protected market is lower than in the unprotected one. Second, the average level of healthiness in the pool of participants in the two markets will in general differ. The fact that a person desires a short-term sexual relationship that does not use a condom signals something about their past tendencies to engage in risky behavior. For instance, in the market for protected sex the fraction of healthy partners $R_p(1)$ might be higher than

the corresponding fraction $R_u(1)$ in the market for unprotected sex. The value function for a healthy individual who engages in sex of type $s = u, p$ is then

$$\begin{aligned}\tilde{V}_s^\beta(1, x) &= \ln(y - z_s) + pI(s) + u[1 - I(s)] \\ &\quad + \sum_{\hat{\phi}} R_s(\hat{\phi})[1 - \gamma_s(\hat{\phi})]\chi(c)\beta[qV_l^\beta(t, x) + (1 - q)V_l^\beta(0, x)] \\ &\quad + \{1 - \sum_{\hat{\phi}} R_s(\hat{\phi})[1 - \gamma_s(\hat{\phi})]\chi(c)\}\beta V_l^\beta(1, x).\end{aligned}\quad (3)$$

It entails the flow utility of consumption and sex and the continuation values. The individual gets infected with probability $\sum_{\hat{\phi}} R_s(\hat{\phi})[1 - \gamma_s(\hat{\phi})]\chi(c)$. This infection probability takes into account the distribution in market s over health status as given by $R_s(\hat{\phi})$. For each health status it factors in the transmission risk, $[1 - \gamma_s(\hat{\phi})]\chi(c)$, which depends on whether the individual is circumcised. If the individual contracts the virus, they have the immediate probability q of getting treated and obtain a treated person's continuation utility, $\beta V_l^\beta(t, x)$, while continuation as an untreated infected individual, which has value $\beta V_l^\beta(0, x)$, arises with probability $1 - q$. With probability $\{1 - \sum_{\hat{\phi}} R_s(\hat{\phi})[1 - \gamma_s(\hat{\phi})]\chi(c)\}$, the individual remains healthy, in which case they obtain the continuation utility associated with being healthy, $\beta V_l^\beta(1, x)$.

Finally, to be matched in the short-term market, a person must first decide how much effort to expend searching in each market; that is, they must choose π_p and π_u . This is done in accordance with the following problem:

$$\begin{aligned}V_s^d(\phi, x) &= \max_{\substack{0 \leq \pi_u^d, \pi_p^d, \\ \pi_u^d + \pi_p^d \leq 1}} \{ \pi_p^d \tilde{V}_p^d(\phi, x) + \pi_u^d \tilde{V}_u^d(\phi, x) + (1 - \pi_p^d - \pi_u^d) \tilde{V}_a^d(\phi, x) \\ &\quad - C(\pi_p^d) - C(\pi_u^d) \}, \text{ for } d = \iota, \beta.\end{aligned}\quad (4)$$

The value function $V_s^d(\phi, x)$ gives the ex-ante value for a type- x individual, with discount factor d and health status ϕ , of entering the market for short-term sex. (Here the subscript s simply denotes short term.) The solution for search effort is represented by the function $\pi_s^d = \Pi_s^d(\phi, x)$, for $s = p, u$.

3.2 Long-term Relationships

Imagine a person with high discount factor, β , who is currently in a long-term relationship. In a long-term relationship there are no choices to make: there are no affairs, all sex is unprotected, and the partnership endures until some form of exogenous breakup occurs. A long-term relationship may end due to an exogenous divorce, the partner dying of natural causes, or the partner developing the symptoms of HIV/AIDS.

To understand the value of a continuing relationship in which an individual of health status ϕ and circumcision type c is married to a partner of status $\hat{\phi}$ and circumcision type \hat{c} , it is important to derive the probability $\Upsilon(\phi', \hat{\phi}' | \phi, \hat{\phi}, c, \hat{c})$ that the pair enters next period together with status ϕ' and $\hat{\phi}'$. A pair who is currently healthy will stay healthy, so that $\Upsilon(1, 1 | 1, 1, c, \hat{c}) = 1$ for all (c, \hat{c}) . To illustrate a more complicated case, the following equation shows the transition to a state next period where the individual is infected and the partner is treated, from a state this period where the individual is healthy and the partner is infected or treated ($\hat{\phi} = 0, t$):

$$\Upsilon(0, t | 1, \hat{\phi}, c, \hat{c}) = [1 - \gamma_u(\hat{\phi})]\chi(c)(1 - q)Q(\hat{\phi}). \quad (5)$$

The first term captures the chance that the individual gets infected $[1 - \gamma_u(\hat{\phi})]\chi(c)$. The second and third terms capture the probabilities that the individual is not treated, $1 - q$, but the partner is, $Q(\hat{\phi})$. All other transition probabilities are described in Appendix D.2.

Given these transitions, consider an individual of health status ϕ who starts the period matched to a partner of health status $\hat{\phi}$ and circumcision type \hat{c} . Focus on a high-discount-factor individual for illustration. His continuation utility is

$$\begin{aligned} \tilde{V}_l^\beta(\phi, \hat{\phi}, \hat{c}, x) &= \ln(y - z_l) + u + l + \alpha_\phi \beta A \\ &+ (1 - \alpha_\phi)(1 - \epsilon)(1 - \delta)(1 - \alpha_{\hat{\phi}})\beta \sum_{\phi', \hat{\phi}'} \Upsilon(\phi', \hat{\phi}' | \phi, \hat{\phi}, c, \hat{c}) \tilde{V}_l^\beta(\phi', \hat{\phi}', \hat{c}, x) \\ &+ (1 - \alpha_\phi)[1 - (1 - \epsilon)(1 - \delta)(1 - \alpha_{\hat{\phi}})]\beta \sum_{\phi', \hat{\phi}'} \Upsilon(\phi', \hat{\phi}' | \phi, \hat{\phi}, c, \hat{c}) V_l^\beta(\phi', x). \end{aligned} \quad (6)$$

The first three terms on the first line capture the flow utility from consumption and sex in long-term relationships. Additionally, with probability α_ϕ the individual develops final stage symptoms which is captured by the last term in the first line. With complementary probability the individual remains either married (second line) or becomes single (third line). The marriage will persist if there is no exogenous breakup (chance $1 - \epsilon$), the partner does not die of natural causes (chance $1 - \delta$) and the partner does not develop symptoms (chance $1 - \alpha_{\hat{\phi}}$). In this case the individual obtains the continuation value connected with marriage, $\beta \tilde{V}_l^\beta(\phi', \hat{\phi}', \hat{c}, x)$, taking into account the transition probability of health status. With complementary probability, the marriage breaks up and the individual carries his new status as a single into the long-term market with associated continuation value, $\beta V_l^\beta(\phi', x)$.

Since individuals do not know their partner's health status, the value of being matched in the long-term market for an individual with discount factor $d = \iota, \beta$ and health status $\phi = 0, 1, t$ is given by the weighted average of possible partners in the long-term market:

$$\tilde{V}_l^d(\phi, x) = \sum_{\hat{\phi}, \hat{c}} R_l(\hat{\phi}, \hat{c}) \tilde{V}_l^d(\phi, \hat{\phi}, \hat{c}, x). \quad (7)$$

Note that, unlike short-term relations, in long-term relationships, the relevant fraction $R_l(\hat{\phi}, \hat{c})$ also depends on the circumcision type of the potential partner. This is so because the circumcision type of the spouse matters for the evolution of the marriage and its duration, as can be seen from the probabilities given in Appendix D.2.²¹

The value of searching in the long-term market, for a status/type- (ϕ, x) person with discount factor d , is given by

$$V_l^d(\phi, x) = \max_{\pi_l^d} [\pi_l^d \tilde{V}_l^d(\phi, x) + (1 - \pi_l^d) V_s^d(\phi, x) - C(\pi_l)], \quad (8)$$

for $d = \iota, \beta$ and $\phi = 0, 1, t$. The solution for search effort, π_l^d , is represented by the function $\pi_l^d = \Pi_l^d(\phi, x)$.

²¹This is only relevant for women as only men get circumcised; i.e., in the model a man cannot have a circumcised mate.

3.3 Stationary Equilibrium

The analysis focuses on a steady-state competitive equilibrium where the number of people of a particular gender, discount factor, permanent type, and treatment status are constant through time. From the above dynamic programming problems, time-invariant decision rules arise for search in each of the three markets. By using these decision rules, the steady-state type distribution can be formulated.²² In the equilibrium modeled, beliefs about the prevalence rates in each market are rational and thereby consistent with the time-invariant decision rules and steady-state type distribution. Decision rules are optimal given prices and beliefs. Finally, prices are market clearing so that the number of men who form a sexual relationship in a given market equals the number of women who form a relationship in that market. Refer to Appendix D.3 for an exact definition of the equilibrium and for more information on its formulation.

4 Calibration

To address the HIV/AIDS epidemic in Malawi, the model is analyzed numerically. A benchmark simulation is constructed that displays features that are broadly consistent with the Malawian case. In particular, the simulated model has an HIV/AIDS infection rate that corresponds with the Malawian data, a proportion of casual sexual encounters that is approximately the same, and a reasonable fraction of these encounters use a condom.

Interpret a model period as lasting one quarter. The model allows for heterogeneity along four dimensions. First, people differ in their discount factor. Second, men and women have different transmission rates, so that the same level of sexual activity leads to a discrepancy in the HIV rate across genders. Third, among men there is one more distinction regarding transmission rates: circumcised men are less likely to contract the virus. Fourth, some infected people receive ART, others don't. Since ART was introduced in Malawi only in 2005, which is in the latter part of the period under study, this dimension of heterogeneity is shut

²²Appendix D.4 outlines how to derive the steady-state type distributions.

down in the benchmark calibration. It will be investigated later in Section 5.

The calibration is conducted in three steps. First, to the extent possible, parameters with direct data analogs are assigned values from the literature. In particular, all parameters relating directly to the biology of the disease are chosen this way. Second, the remaining parameters are selected to match some key observations related to the HIV/AIDS epidemic in Malawi. The data mostly obtains from the 2004 Demographic and Health Survey (DHS) that was conducted in Malawi. Using the micro data from this survey, a number of statistics are computed regarding HIV prevalence rates, sexual behavior, marital status, etc. Third, the model's predictions are compared to the Malawian data along several non-targeted dimensions. The model performs surprisingly well, which can be interpreted as an additional validity check.

4.1 Parameters Based on Direct Evidence

The most important parameter values for the simulation are those concerning HIV/AIDS. Fortunately, for the most part, these can be taken from the medical literature. The transmission risk for one-time male unprotected sex is taken to be 4.5 per 1,000 encounters. This number falls in the range of estimates reported by a variety of studies.²³ Since couples on average have sex 9 times a month, as reported in Gray et al. (2001), this translates into a quarterly non-transmission risk of $\gamma_u^m = 0.879$.²⁴ The transmission risk when condoms are used is obviously lower, but protection is far from perfect—Bracher, Santow, and Watkins (2004). Select $\gamma_p^m = 0.96$, corresponding to a 67% efficacy rate, which is in line with Weller (1993) who conducted a meta-analysis of condom efficacy. Assume that circumcised men are 60% less likely to contract the HIV/AIDS virus and set $\chi = 0.4$ accordingly. This is in accord with the improvements reported by Auvert et al.

²³For example, the annual report of the UNAIDS Joint United Programme on HIV/AIDS (2007) gives a range of 2 to 6 per 1,000 encounters, depending on whether other STDs are present or not. Baeten et al. (2005) also report a transmission risk of 6 per 1,000. Gray et al. (2001) report a somewhat lower number of 1.1 per 1,000 for Uganda; however, free condoms were distributed as part of the study. Wawer et al. (2005) finds transmission rates as high as 82 per 1,000 encounters during the first few months after infection.

²⁴This quarterly rate also is similar in magnitude to the per-partnership transmission rates for Sub-Saharan Africa reported in Oster (2005) and the correction appendix.

(2005), Bailey et al. (2007) and Gray et al. (2007). Set the fraction of circumcised men to 20% (DHS 2004). Further, for physiological and anatomical reasons, and in accord with the medical evidence, women are assumed to have a higher risk of contracting HIV than men. Nicolosi et al. (1994) reports a risk that is 2.3 times as high for women.²⁵ However, the range of estimates is wide. At one extreme, Gray et al. (2001) find no statistically significant difference between transmission rates by gender. At the other extreme, Padian, Shiboski, and Jewell (1991) calculate a factor as high as 20. Assuming that women are 75% more likely to get infected implies a one-time risk of 7.87 per 1,000 for unprotected sex. This delivers a quarterly non-transmission rate of $\gamma_u^f = 0.787$. Using the same gender gap for protected sex yields $\gamma_p^f = 0.929$.

The average time from infection to the outbreak of symptoms is equal to 10 years (DHS Report, 2004). Therefore, let $\alpha = 0.025$; i.e., 40 quarters. The average time from the outbreak of symptoms to death is 2 years (DHS Report, 2004). Thus, pick $\delta_2 = 0.125$; i.e., 8 quarters.

Some other parameter values can also be pinned down using a priori information. Set the quarterly divorce hazard equal to $\epsilon = 0.03$. Bracher, Santow, and Watkins (2004) report that 26.4% of all marriages in Malawi end in divorce within the first five years. Assuming a constant annual divorce hazard, this would imply a quarterly risk of 1.56%.²⁶ A rate twice this number is used: First, polygyny is fairly common in Malawi, from which the analysis abstracts. Second, extra-marital affairs are relatively common as well. Therefore, interpret, for example, a 10-year marriage with one affair as two long-term relationships with a third casual one in between.²⁷

The quarterly (non-HIV related) death hazard is taken to be $\delta = 0.006$. A study conducted by the U.S. Census Bureau (2004) reports a life expectancy without HIV of 56.3 years for Malawi. Since the model starts at age 15, the quarterly death hazard is selected to match a life expectancy of 41.3 years. Malawi is a very poor country. Set $y = \$320$ (U.S.), which corresponds roughly to quarterly GDP

²⁵Oster (2005) also reports a factor of two in her measures of per-partnership transmission rates.

²⁶A similar number is also reported by Reniers (2003).

²⁷Greenwood et al. (2017b) explore what happens to equilibrium outcomes when the risk of divorce is lower.

per working age population in 2001 (note that only about half the population is of working age in Malawi). Table 1 summarizes the preceding paragraphs by listing all parameters that are set a priori.

TABLE 1: PARAMETERS CHOSEN OUTSIDE THE MODEL

<i>Parameter</i>	<i>Value</i>	<i>Interpretation</i>
γ_u^m	0.879	12.1% quarterly transmission risk, unprotected sex, uncircumcised men
γ_p^m	0.96	4% quarterly transmission risk, protected sex, uncircumcised men
χ	0.4	circumcised men are 60% less likely to contract the virus
γ_u^f	0.787	21.3% quarterly transmission risk, unprotected sex, women
γ_p^f	0.929	7.1% quarterly transmission risk, protected sex, women
α	0.025	10 years from infection to symptoms
δ	0.006	6% quarterly death risk
δ_2	0.125	2 years from symptoms to death
ϵ	0.03	3% quarterly divorce hazard
y	320	quarterly income

4.2 Parameters Chosen to Match Data Moments

The remaining parameters have no clear data analogs. For example, the utilities from the different types of sexual relationships are free parameters, constrained only by $p < u$; i.e., people enjoy unprotected sex more than protected sex. The values for these parameters are picked to match several facts related to sex, marriage, and HIV/AIDS in Malawi.

As specified above, the only exogenous heterogeneity (in addition to the difference in transmission risks for circumcised and uncircumcised males and the difference in transmission risks across genders) is the degree of patience people have. Recall that $\tilde{\iota}$ and $\tilde{\beta}$ denote the discount factors for the young and the old. Note that these are “pure” discount factors; i.e., net of mortality risk.²⁸ Suppose that $\tilde{\beta}$ varies across individuals according to a uniform distribution with support $[\tilde{\beta}_{min}, \tilde{\beta}_{max}]$. Moreover, assume the ratio of discount factors when young and old is always given by the same value $\iota_{change} = \tilde{\iota}/\tilde{\beta} < 1$. Thus, there are 11 free parameters: $p, u, \ell, \omega_s, \omega_l, \kappa, A, \eta, \tilde{\beta}_{min}, \tilde{\beta}_{max}$, and ι_{change} . See Table 2 for a summary.

²⁸That is, $\beta \equiv \tilde{\beta}(1 - \delta)$ and $\iota \equiv \tilde{\iota}(1 - \delta)$.

TABLE 2: CALIBRATED PARAMETERS

<i>Interpretation</i>	<i>Parameter value</i>
Flow utility unprotected sex	$u = 7.8$
Flow utility protected sex	$p = 1.4$
Flow utility long-term sex	$l = -4.8$
Discount factor, min and max support	$\tilde{\beta}_{min} = 0.969, \tilde{\beta}_{max} = 0.9999$
Ratio discount factors, young vs. old	$\iota_{change} = 0.874$
Value of life with Aids	$A = 5.8$
Prob. of switch to high discount factor	$\eta = 0.116$
Search cost parameters	$\omega_s = 0.44, \omega_l = 17.5, \kappa = 0.115$

To discipline the choice of the parameters, 11 data moments are targeted. Table 3 gives an overview of the data moments and shows how well the benchmark model matches them.²⁹

The main targets are the overall prevalence rate for HIV/AIDS in society, and the prevalence rates for each gender. Some additional targets are selected to discipline the exercise, such as the fraction of sex that is casual as opposed to long-term, the fraction of the population that is single, and the pattern of marriage by age. This ensures that there is not too much reliance on risky short-term interactions. Many people are married, and get married quickly. Moreover, women get married faster. Additionally, the fraction of condom users in casual encounters and the number of such encounters are considered. This is important because the model emphasizes the choice of whether to engage in risky sexual activity. Finally, the fraction of people that die of natural causes, as opposed to HIV/AIDS, is included for overall consistency.

The upshot of the analysis is that the benchmark simulation matches these key features of the Malawian HIV/AIDS epidemic pretty well. To begin with, as can be seen from Table 3, the HIV/AIDS prevalence rate predicted by the model is 10.3%, close to the 11.8% in the data. Moreover, the benchmark simulation captures the gender difference in HIV/AIDS infection rates, with females experiencing a rate that is 3.5 percentage points higher than that for males (12.1% versus 8.6%). The model also captures well the HIV incidence rate (which was not specifically targeted). In the model 11.7 per 1,000 healthy people get newly

²⁹All data sources for the tables and figures are discussed in Appendix B.

TABLE 3: TARGETED MOMENTS

<i>Observation</i>	<i>Data</i>	<i>Model</i>
HIV/AIDS prevalence rate, %	11.8	10.3
—Males	10	8.6
—Females	13	12.1
Sex that is casual, % (of all)	18	16
Condom use for casual sex, %	39	33
Singles that had casual sex in past year, %	47	54
Singles, %	33	48
Married by age 22, %		
—Males	58	57
—Females	90	63
Married by age 50, %		
—Males	100	98
—Females	100	98
Deaths related to HIV, %	29	25

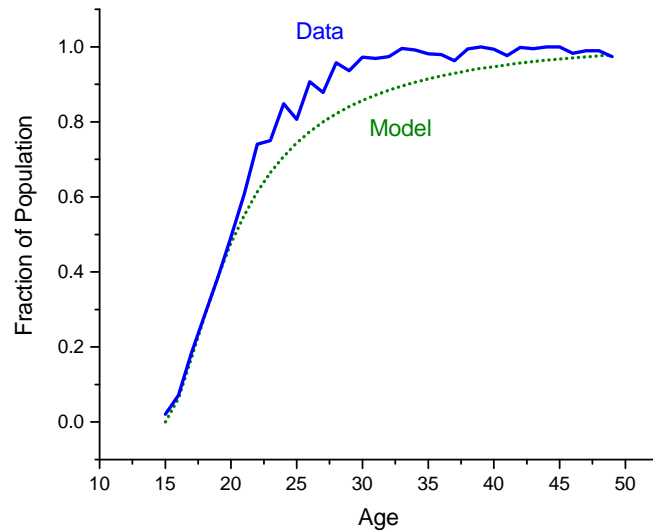
infected every year, compared to 11.1 in the data.

In addition to matching moments on HIV/AIDS prevalence, the benchmark model also matches moments of other aspects of sexual activity quite well. In the model casual or short-term sex is a small fraction of all sexual encounters: 16%, close to the 18% of sex that occurs outside of a union that is reported in the data. For those who engage in casual sex, the model predicts that 33% use a condom. This is less than the 39% seen in the data, but is still close. In fact, as people have been found to overstate the amount of protected sex they have [see Allen et al. (2003)], these two numbers may be closer in reality than first meets the eye. The next row in Table 3 reports the fraction of singles who had casual sex in the last year. These statistics are different from the fraction of all sexual activity that is casual because (all) married people have sex while some singles are abstinent. Singles in the model have a little more casual sex than their real-life counterparts (54 versus 47%). Finally, the fraction of the population that dies from HIV/AIDS is comparable across model and data (25 versus 29%).³⁰

Table 3 also lists the fraction of singles in the entire population. The model pre-

³⁰Appendix C provides sensitivity analysis for these estimated parameters. Both benchmark moments and the response to some policy experiments (discussed in Section 5) are analyzed.

Figure 1: Fraction of Population ever Married—Model vs. Data



dicts this fraction to be 48%, somewhat higher than the 33% observed in the data, which is partially due to the choice to break up marriages more frequently in the model than in the data to account for extramarital affairs. The model captures some of the gender differences in the timing of marriage. Women marry much earlier than men—in the data, 90% of women are married by age 22, whereas only 58% of men are married by this age. The corresponding numbers in the model are 63% and 57%. The model generates the earlier marriage of women (relative to men) via their higher infection risk. This makes the safety of marriage more attractive for women vis à vis men. In reality HIV risk is only one reason why women get married earlier than men. Additional considerations such as the risk of pregnancy make casual encounters more risky for women, which might explain the larger wedge between the genders in the data compared to the model. Polygyny may contribute to the age difference in marriage since men must save in order to be marriageable because potential wives are now scarcer. As is only reasonable, men eventually “catch up,” and by age 50 almost everyone is married, both in the model and the data. See Figure 1 for a comparison of the fraction of the population that has ever married in the model vs. the data.

The benchmark model generates some other predictions that are not targeted

when picking values for parameters. Protected sex in the benchmark simulation is substantially cheaper for men than unprotected sex ($z_p = -6.5 < z_u = 277$). Such a premium has indeed been documented in the literature. Based on data from commercial sex workers in Mexico, Gertler, Shah, and Bertozzi (2005) report a 23% premium for unprotected sex. The premium increases to 46% when the sex worker is considered to be very attractive. Further, Appendix E compares several life-cycle implications of the model to the data. The model tracks these data patterns surprisingly well despite the limited degree of heterogeneity and despite the limited state variables that describe the agent's life-cycle. In Section 5 the model's predictions will be compared with cross-country data on the relationship between HIV and male circumcision. Moreover, in Section 6, time series data for Malawi will be confronted with a version of the model with a richer information structure regarding knowledge about HIV. This all provides additional validity checks for the model.

5 Policy Experiments

The model is now ready to explore the effectiveness of various policies intended to curb the spread of HIV/AIDS. This section investigates four specific policies. The first two policies have been at the forefront of the policy discussion—namely the introduction of antiretroviral therapy (ART) and circumcising males. The two additional policies yield some surprising results—improving condoms and treating other sexually transmitted diseases. A companion paper analyzes further policies aimed at the behavioral aspects of the problem, namely policies aimed at increasing marriage and thereby reducing casual sex (Greenwood et al. 2017b).

In addition to studying the effectiveness of the various policies in the full model, two alternative versions of the model are simulated: (i) small-scale field experiments and (ii) epidemiological experiments. To be concrete, the field experiment assumes that only a small fraction of the population is treated and changes their behavior, but that this fraction interacts in equilibrium with everyone else at the *preexisting* equilibrium prices and infection rates. In other words the field experiment is a partial equilibrium version of the general equilibrium one. The epi-

demological experiment assumes that people make no behavioral adjustments. It therefore uses the policy functions from the benchmark calibration. The infection probabilities and assessments of transmission risks are governed by the new transmission probabilities. Comparing the effectiveness of policies across the three versions of the model can be thought of in two ways. First, they can be viewed of as thought experiments. A comparison of the small-scale field experiment with the benchmark model experiment shows the importance of general equilibrium effects, while a comparison of the epidemiological experiment with the benchmark model illustrates the significance of behavioral responses. Second, they give insights into the extent to which actual field experiments and epidemiological studies might (or might not) generate reliable policy advice. As will be seen below, results from the full equilibrium experiment are often in between the epidemiological and the field experiments. However, depending on the type of policy, there are exceptions.

In this context it must also be stressed that there are multiple types of equilibrium effects. First, the aggregate prevalence rate changes, which changes the riskiness of sex, and thereby can amplify (or mitigate) the direct impact of a policy. Second, there is also a secondary behavioral adjustment in response to the change in the aggregate prevalence rate. Third, prices change and people may change their behavior in response. Note that while equilibrium effects through prices are generally front and central in economics, these price effects are the least important in the current context. Most importantly, behavioral adjustments in response to price changes are very subtle as men and women are affected in opposite ways. Assume, for example, the price for sex with condoms goes up. Would people use fewer condoms? This is not obvious as an increase in the price means that men will have to pay more for protected sex, but also that women receive more. Thus, all else equal, men would demand less protected sex, but women would demand more. Since this cannot be an equilibrium, something else needs to adjust to assure market clearing. Thus, the aggregate amount of a particular type of sex is not necessarily monotone in the price.

Table 4: HIV AND CIRCUMCISION ACROSS COUNTRIES—REGRESSIONS

	Dep. variable: HIV prevalence rate				Dep. variable: HIV incidence rate			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
circumcision	-0.1122***	-0.0765**	-0.0796**	-0.064	-9.840**	-4.972	-6.191	-7.339
Log GDP p.c.	0.0314***	0.0293***	0.0288***	0.0296***	3.87***	3.73***	3.43***	2.459**
ART	0.0816	0.104**	0.105*	0.098	5.63	8.71	9.05	6.266
syphilis	0.0025	0.0029	0.003	0.0045	0.359	0.42	0.526	0.711
Muslim		-0.002	-0.00056	-0.0012		-0.026	-0.128	-0.207*
Christian			-0.00039	-0.00065			-0.121	-0.171*
condom price				-0.268*				-17.5
R^2	0.72	0.73	0.74	0.79	0.61	0.65	0.67	0.71
N	32	31	31	23	30	29	29	22

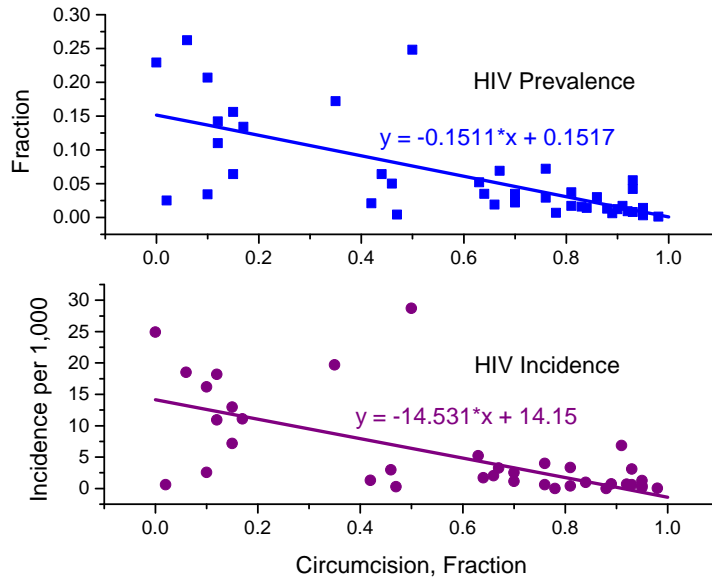
5.1 Male Circumcision and HIV

A policy intervention that has received a lot of recent attention is male circumcision. UNAIDS lists circumcision as one of five prevention pillars in their 2016 report (UNAIDS 2016). Similarly, the World Health Organization has been promoting voluntary male circumcision as a prevention tool (World Health Organization 2016). Consider analyzing this particular policy within the model. In what follows an increase in the circumcision rate from the 20% level in the benchmark model is analyzed.

Start by asking to what extent cross-country variation in circumcision explains the observed differences in HIV prevalence and incidence rates across countries. While circumcision has been advocated recently as a medical policy, it has not been implemented on a large scale yet. Thus, existing cross-country differences in circumcision rates are unrelated to HIV, but rather due to cultural reasons. Data from 32 sub-Saharan African countries is used to explore the empirical relationship between HIV and circumcision.³¹ As Figure 2 shows, there is a clear negative correlation between circumcision and HIV prevalence/incidence in the data. Obviously countries differ along many dimensions other than circumcision that also affect HIV rates, such as how rich a country is. Table 4 reports a more comprehensive empirical analysis, controlling for various potentially confounding factors such as GDP, ART, religion, and the price of condoms.

³¹See Appendix B for details on the data.

Figure 2: Circumcision and HIV across Countries



Now contemplate performing the analogous exercise in the model. In other words, vary the fraction of men who are circumcised in the model and compute the implied equilibrium HIV prevalence and incidence rates. Start with prevalence. There is an almost linear relationship with a slope of -0.05. In other words, for each 10 percentage point increase in circumcision, the HIV prevalence rate declines by about half a percent. This is quite similar in magnitude to the coefficient from the regressions. Controlling for observables, regression (4) in Table 4 shows a coefficient of -0.064. Since the prevalence rate is a stock variable, another way is to compare the incidence of HIV. In the model, the relationship between HIV incidence and circumcision rates yields a coefficient of -6.389. Looking at regressions (5) to (8) in Table 4, the coefficients estimated in the data are close to the model counterpart. Thus, the model quite closely replicates the cross-country relationship between HIV and circumcision. Since cross-country data was not used in the calibration, this can be seen as an external validation of the model.

So is male circumcision a promising policy that should be promoted? Probably yes.³² Table 5 shows that scaling circumcision up from the current 20% to 100%

³²Of course this analysis ignores any potential negative physiological and psychological side effects.

Table 5: CIRCUMCISION

	BM (20% circ)	BM by type		100% circ	Epidem.
		not circ	circ		
HIV prevalence, %	10.3			5.6	4.3
—Males	8.6	8.75	8.0	3.8	4.1
—Females	12.1			7.6	4.6
Sex that is casual (males), %	16	14	22	29	—
Condom use for casual sex (males), %	33	35	27	22	—
Single men, %	50	49	53	59	—
Casual sex in past year, single men, %	21	19	28	31	—
Price—protected	-6.5	—	—	53	—
Price—unprotected	278	—	—	309	—
Price—long term	125	—	—	161	—

would cut the HIV rate by almost half—from 10.3 to 5.6%. This would work despite significant behavioral adjustments by men. The model suggests that men would engage in more risky behavior along all dimensions: less marriage, decreased condom use, and more sex.³³ The equilibrium model also shows a sizable positive effect for women: The prevalence rate among women falls from 12.1 to 7.6%. This is worth emphasizing as the effect on women is theoretically unclear. Women do not directly benefit from circumcision as the male-to-female transmission rate does not change. Moreover, the increased risky behavior of their partners would, all else equal, lead to more HIV among women. Yet, women benefit through the equilibrium effects. Thus, the lower female-to-male transmission rate eventually also leads to a decline in the female prevalence rate.

The conclusion that circumcision is an effective policy is in line with findings in field experiments. Padian et al. (2010) compare the effects of 36 randomized controlled trials (RCTs) in the context of HIV. Only six RCTs delivered definitive results on HIV—three of which were circumcision RCTs. But what about the size of the effect? In the model, the behavior of circumcised men corresponds to the treated group in a small field experiment, while non-circumcised men would constitute the control group. Clearly, in line with RCTs, the treated men in the model have a lower HIV rate. However, the effect is small (8 versus 8.75%). This is due to the missing equilibrium effects and behavioral responses among

³³This is in line with cross-sectional studies showing that circumcised men engage in more risky sexual behavior along several dimensions (Bailey and Othieno 1999; Seed et al. 1995).

the circumcised that crowd out some of the gains from the intervention. The simulation suggests that circumcised men have 50% more casual sex, use condoms about 25% less, and are slightly more likely to be unmarried. The result that circumcised men engage in more risky behavior is in line with the empirical findings from RCTs. Bailey et al. (2007) find that unprotected sex fell a lot less for the newly circumcised relative to the control group.³⁴ Similarly, Auvert et al. (2005) find a statistically significant increase in the number of partners for the circumcised men. In a unique long-run follow-up after a circumcision RCT with teenagers, Kim et al. (2018) find evidence of risk compensation four years after the intervention. Those treated (i.e., being offered free circumcision) have a higher incidence of other STDs and are less likely to believe that they should use a condom during a casual sexual encounter.

A simple extrapolation from such a field experiment might suggest that circumcision is useful but effects are expected to be small. Yet, the experiments show when circumcising the entire male population powerful equilibrium effects kick in. The male HIV rate falls by more than half (from 8.6 to 3.8%). This is despite even larger behavioral responses in equilibrium. The relevant equilibrium effect here is perhaps best compared to interest rate compounding. Because female-to-male transmission rates are lower, for a given sexual behavior, fewer men get infected. Given that fewer men are infected, for a given female sexual behavior, fewer women get infected. And so on.

The equilibrium effect would of course also be present in epidemiological studies. The epidemiological model version is reported in the last column of the table. Not surprisingly, the epidemiological experiment exaggerates the beneficial impact of the policy as it completely ignores the additional risk-taking behavior. The potential overstatement of effects is a concern of the epidemiological literature. Williams et al. (2006) simulate large positive effects of circumcision in Sub-Saharan Africa, but add a cautionary note at the end noting that increases in risk-taking may reduce some of the benefits of male circumcision (MC). The

³⁴Note that most circumcision RCTs were combined with counseling sessions for both the treatment and control group. Since this increased awareness likely led to more cautious sexual behavior, the relevant thing to look at is behavioral change in the treatment relative to the control group, rather than the absolute change. Section 6 discusses the impact of increased awareness about the HIV transmission mechanism.

study concludes “population level studies of MC are now needed to determine the likelihood of behavioral disinhibition and to assess its impact on transmission in the long term.” A similar call for male circumcision studies at scale is made in De Walque (2012). Since population level studies are expensive and difficult to implement, the equilibrium model offers a promising alternative. It appears that the equilibrium effects are much more powerful than the (adverse) behavioral responses, and thus the true beneficial effects of circumcising all men is much closer to what an epidemiological study would find than to a direct extrapolation from a field experiment.

5.2 Was ART successful in reducing HIV?

A second policy that has gained widespread attention is the introduction of antiretroviral therapy (ART). While initially invented as a treatment for sick people, it is now also believed to have a preventive component. ART lowers the viral load and thus makes the person taking the drugs both feel healthier (plus live longer) and less likely to pass the virus on. Since the existence of ART makes life with HIV more tolerable, this may lead healthy people to engage in riskier behavior. Moreover, since HIV-infected people on ART live longer, they have more time to pass the virus on. Thus, the net effect of ART on HIV is not obvious. Previous research, based on different methodologies, finds a wide range of effects. The predicted long-run effects range from the complete eradication of HIV to an increase in the prevalence rate. Much of the medical literature seems convinced of the effectiveness of the policy (e.g. Cohen, M.S. et al (2011)), while for example Lakdawalla, Sood, and Goldman (2006) show empirically that ART has led to an increase in HIV in the United States. Wilson (2012) provides a survey of the empirical studies to date and warns that widespread enthusiasm for treatment as prevention may be misguided, since expected outcomes are currently mostly based on clinical trials alone that are not informative for what would happen if the entire population was treated.

In Malawi, ART was introduced in 2005. Figure 3 shows that over the course of 2004 to 2014, the fraction of infected people on ART increased from essentially

zero to 50%. At the same time, HIV has been declining.³⁵ The Malawian government seems to have concluded that the decline was due to the successful ART scale-up.³⁶ Clearly, by just inspecting the two time series in Figure 3, it is clear that ART cannot be the whole story. The HIV decline started in 2000, which is 5 years prior to the introduction of ART. Anticipation effects would go in the wrong direction—as anticipating ART should make people behave in a more risky fashion without yet experiencing the benefits of the lower transmission risk. It is of course possible that the introduction of ART did contribute to the HIV decline in the later years, however. The model can be used to assess (and quantify) this hypothesis.

ART is modeled as a decline in the (out-going) transmission rate. Assume a reduction in infectivity by two thirds, which is within the range of empirical estimates.³⁷ Specifically the quarterly transmission risk for unprotected sex declines from 0.21 to 0.07 for females having sex with someone on ART, and from 0.12 to 0.04 for males.³⁸ The reduced mortality (and accordingly increased quality of life) is modeled as a longer life without symptoms. This means that infected people on ART live longer, but also enjoy a better life after infection (relative to those not on drugs). Specifically reduce α , the probability of symptoms breaking out, from 0.025 to 0.0125. Since symptoms are quickly followed by death, this means that mortality is essentially reduced by 50% which is in line with the evidence.³⁹

³⁵The time series data for the HIV prevalence rate comes from UNAIDS and is estimated largely based on information collected at antenatal clinics. The UNAIDS 2004 HIV rate is 13.5% which is substantially higher than the prevalence rate in the DHS, reported in Table 3. Since the DHS is a nationally representative survey, it is the more reliable source. Yet, UNAIDS provides the only consistent time series data.

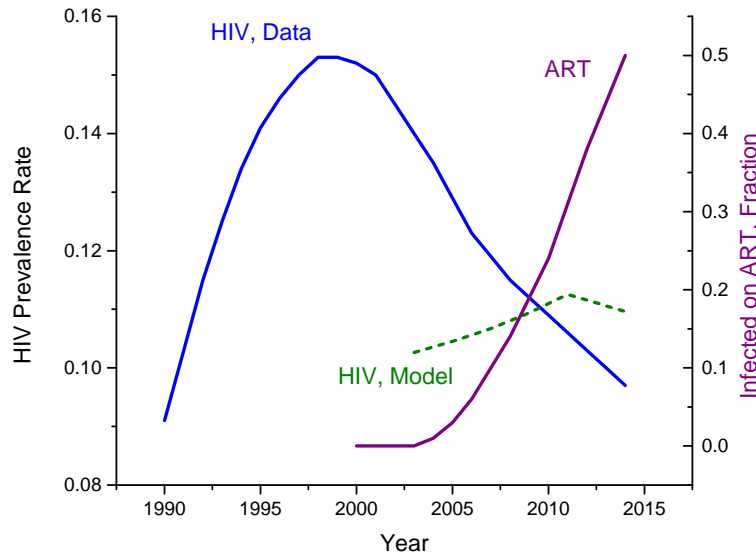
³⁶“Malawi’s rapid and successful Antiretroviral Therapy scale-up from 2004 to 2014 has critically influenced the trajectory of the HIV epidemic ...”, see p.2 in the Malawi AIDS Response Progress Report 2015, Government of Malawi.

³⁷The estimates cover a wide spectrum, ranging from a low reduction of 13% in Lakdawalla, Sood, and Goldman (2006), to a medium decline of 60% in Porco et al. (2004) to a high reduction of 94% in Cohen, M.S. et al (2011). Most of the high numbers are related to the best possible treatment; i.e., early introduction of a combination antiretroviral therapy, which is not the standard treatment in poor countries today. The timing seems important for the reduction of infectivity, with a higher reduction when treatment is given right after infection, something the model does not capture. Thus, for the purposes here, the more conservative estimates seem more relevant.

³⁸This corresponds to an increase in γ from 0.787 to 0.929 for females and from 0.879 to 0.96 males. The numbers for protected sex and for circumcised men are adjusted accordingly when having sex with a treated individual.

³⁹A study specifically for Malawi is Lowrance (2007). Group (2015) finds similar mortality

Figure 3: HIV and ART in Malawi: Model vs. Data

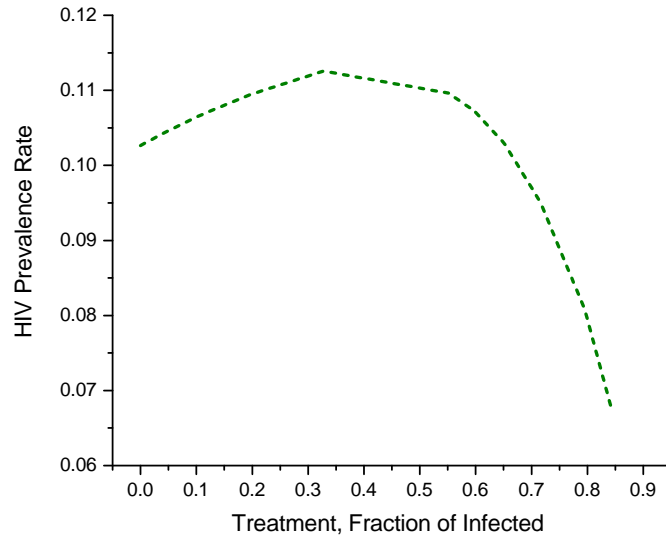


In the model, infected people are treated randomly. The probability of being selected for treatment is q . Since treatment is an absorbing state, $q\%$ newly treated each period cumulates to a substantially higher ART rate. Envisage the following exercise: increase q such that the equilibrium fraction of infected on ART goes up in line with the data. The model then gives, at various levels of treatment, the long-term HIV rate. Since the model compares steady states (i.e., ignores transitional dynamics), this exercise will give an upper bound on the fraction of the HIV decline due to ART.

The dashed line in Figure 3 displays the simulation results. The surprising answer is that hardly any of the HIV decline can be attributed to the introduction of ART. In other words, the negative effects (increased risk-taking, longer time to infect others) of the policy dominate the positive effects (lower transmission rate making sex safer). Now, should one conclude from this that ART is not an effective policy to curb HIV? Probably not. The experiments show that the relative importance of the positive versus negative effects changes with the ART rate. Figure 4 shows that higher levels of treatment are actually quite effective.

reductions for the Ivory coast and Ray (2010) for the United States and Europe. Friedman (2015) also feeds in a 50% reduction in mortality for treated individuals in her simulation study.

Figure 4: ART in the Model



Once more than 50% of the infected are treated, the preventive effect dominates. In other words, going forward, a further expansion of the ART program appears to be a promising idea. The point that ART may be beneficial only once a large enough fraction of the infected is treated is also emphasized in Friedman (2015), who simulates a one-gender model with only one margin of risky behavior.

To better understand the nonmonotonicity, also consult Table 6. When a third of the infected is treated, people realize they may have a reasonable chance of getting treated, if they catch HIV and accordingly take less precautions. Condom use goes down by six percentage points, casual sex goes up from 54 to 63%, and there are more singles now. When 80% of the infected are treated, the behavioral adjustments are even more dramatic. The increase in risk-taking is in line with the empirical findings. Most evidence to date comes from developed countries.⁴⁰ However, there is some limited evidence from African countries as well. Identifying the behavioral response to ART is notoriously difficult. Clearly a randomized controlled trial to assess the effect on the uninfected is difficult to

⁴⁰For example, see Crepaz, Hart, and Marks (2004) for a meta-analysis of 25 studies in the United States, Europe, and Australia.

conceive.⁴¹ Some studies use heterogeneity in beliefs about the effectiveness of ART, such as De Walque, Kazianga, and Over (2012) who find a large increase in self-reported risky sexual activity in response to ART in Mozambique—both in terms of more casual sex and less condom use. Similarly, Cohen et al. (2009) find increased risk-taking behavior in those men who have stronger beliefs about the effectiveness of ART compared to those with more skeptical beliefs. Friedman (2015) uses a difference-in-difference strategy based on proximity to an ART provider in Kenya. She finds an increase in self-reported risky behavior of about 40%—both in terms of the incidence of casual sex and condom use. Using a biomarker (pregnancy) the effect is even larger, about 80%.⁴²

Behavioral adjustments alone would lead to a large increase in the HIV rate, as the field experiment shows, which predicts an increase in the HIV rate to 16.7%. However, these effects are mitigated in equilibrium by the fact that the treated interact with everyone else. Since in the full general equilibrium experiment 80% of the population is treated, sex in general is safer, which leads to a lower aggregate prevalence rate. In the field experiment, people realize that once they get sick they have a 80% chance of getting treated eventually (7.5% per period) but they still have sex with the general population; i.e., essentially none of their potential partners is treated. So their behavioral adjustments lead to a very high HIV rate of 16.7% and there is no offsetting effect through a lower chance of catching the virus. By contrast, by completely ignoring the behavioral adjustments in the epidemiological experiment, one would make the opposite mistake. The epidemiological version predicts a massive HIV decline to a prevalence rate of only 2.5%. In other words, the behavior and the equilibrium effects go in opposite directions—which of them dominates depends on the fraction of people treated. A too cautious attempt may backfire.

Finally, note that this section referred to HIV prevalence, but the same insights apply to incidence, which fell even more strongly over time as it is a flow variable.

⁴¹In the model field experiment it is easy to “treat” some uninfected by changing their q ; i.e., the probability that they would get treated with ART should they get sick. Implementing this in reality would require essentially an information treatment, where some people are informed about increased access to ART while others are not.

⁴²Friedman (2015) provides a discussion on why self-reported sexual behavior changes may be biased downward. Hence, biomarkers may give the more reliable picture.

Table 6: ART

	<i>Benchmark</i>	<i>q=0.01</i>		<i>q=0.075</i>	
		G. E.	G. E.	Small Field	Epidem.
Fraction on ART	none	33	80	80	80
HIV prevalence, %	10.3	11.3	8.1	16.7	2.5
Casual sex, % (of all)	15.7	20.5	42	22	—
Casual sex with condom, %	33	27	24	30	—
Singles who have casual sex, %	54	63	84	62	—
Single men, %	50	53	67	54	—
Single women, %	46	49	65	51	—
Price—protected	-6.5	7.8	25	—	—
Price—unprotected	279	285	259	—	—
Price—long term	125	126	100	—	—

Since the model is in steady state, it induces a tight, increasing and nearly-linear relationship between prevalence and incidence, and therefore insights from prevalence carry over: in the model there is no effect of ART on incidence before 2005, afterwards incidence first rises before declining at modest scale, and only large-scale treatment would induce large changes in incidence. Because of the tight relationship between prevalence and incidence some model implications in the following sections are explicated only for one of them — usually for prevalence.

5.3 Other Policies

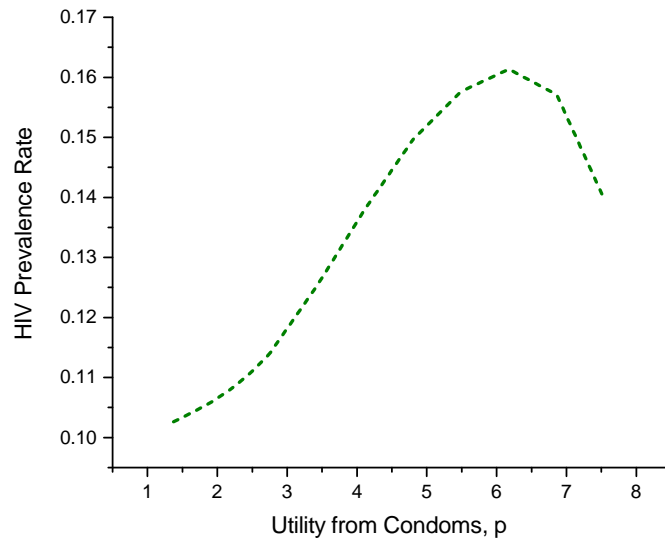
Consider now two additional policies. The first is about condoms and the second about treating other sexually transmitted diseases.

5.3.1 Better Condoms

Suppose one could design more pleasurable condoms (or perhaps raise the psychic pleasure of sex with a condom by reducing stigmatization through publicity campaigns).⁴³ Would this be desirable? The effect of more pleasurable condoms

⁴³While this might seem somewhat far-fetched, note that UNAIDS lists exactly such a policy in their recent report: “Develop new approaches to increase condom use and to enhance the positive perception of condoms among the various populations in need,” p. 30 in UNAIDS (2016). Concerns about the acceptability of condoms have long been discussed in the literature — see the

Figure 5: Better Condoms



is displayed in Figure 5 where starting from the benchmark, utility from sex with condoms is increased until it reaches the same period utility as unprotected sex. It turns out that the HIV rate displays a nonmonotonic pattern when increasing pleasure from condoms, p . The reason that increasing the utility from protected sex does not always lead to a lower prevalence rate is that single life becomes more attractive. So, even though condom usage increases, there are more singles in total and they engage in both protected and unprotected sex. Table 7 gives the example of quadrupling condom pleasure (from 1.4 to 5.5) and then increasing it even further, to 7.5 at which point protected and unprotected sex give almost equal utility. In response, condom use increases tremendously, almost doubling from 32 to 59%. However, the fraction of single men and women also substantially increases (from 48 to 62%). Moreover, the fraction of singles who engage in short-term sex goes up from 54 to 66%. In sum, while more people use condoms, there is also more casual sex. These two forces (safer sex versus more sex) push the prevalence rate in opposite directions. In the example of quadrupling the pleasure, the latter force dominates, so that the overall HIV rate goes up by about 60% (from 10 to 16%). Yet, that is not always the case as can be seen from introduction and Footnote 4.

Table 7: BETTER CONDOMS

	Benchmark	G. E.	Small Field	G. E.	Small Field
p	1.4	<i>Better</i> (5.5)		<i>Better still</i> (7.5)	
p/u	0.18	0.70		0.97	
HIV prevalence, %	10.3	15.8	15.6	14.0	18.8
Casual sex, % (of all)	15.7	34	31	32	41
Casual sex with condom, %	33.0	59	57	62	61
Singles who have casual sex, %	54.0	66	73	73	78
Single men, %	50	64	60	62	70
Single women,	46	60	51	58	56
Price—protected	-6.5	246	–	260	–
Price—unprotected	279	264	–	244	–
Price—long term	125	134	–	138	–

Figure 5. As condoms first get better, the HIV/AIDS prevalence rate initially goes up. As more and more people start to use condoms, however, the lower likelihood of getting infected becomes more important and the prevalence rate starts to go back down. This experiment highlights the potential of some policies to backfire and actually *increase* the overall prevalence rate. These effects can be quantitatively quite important: in the experiment, the HIV rate goes up by 40% as the utility gap between protected and unprotected sex disappears.

Implementing this particular policy as a field experiment gives surprising results. Depending on the exact increase in pleasure, the field experiment effects are larger or smaller than the full equilibrium effects—see Table 7. This is due to two competing forces. Recall that several types of equilibrium effects are not present in the field experiment. In particular the aggregate prevalence rate does not adjust, and hence there are no behavioral responses to the change in prevalence rate either. As the aggregate prevalence rate goes up in this experiment, people realize that sex is more risky and hence use even more condoms. More generally, as the table shows, the behavioral response in risky behavior is more pronounced in the field experiment as singles have more casual sex.⁴⁴ In sum, there are two forces pointing in opposite directions. On the one hand, the behavioral response is larger in the field experiment, which dominates for the second

⁴⁴Marriage goes down by *more* in equilibrium than the field experiment, which may seem puzzling at first. Note, however, that sex in marriage is unprotected by assumption. So marriage itself now is a more risky activity. This effect again is absent in the field experiment version.

experiment (increasing condom pleasure to 7.5). On the other hand, the compounding coming from the aggregate prevalence rate increasing and hence amplifying any given behavior change is only present in the full equilibrium model, which dominates in the first experiment (5.5). Thus, the reaction in the field experiment can be amplified or mitigated relative to the full equilibrium effect.

Finally, an epidemiological experiment would predict no effect of the condom policy. This is by construction, as epidemiological experiments assume no change in behavior, but without behavioral change, the increased condom pleasure by itself would not do anything. This is worth pointing out, since in the case of the medical policies the lack of behavioral adjustments leads to an exaggeration of effects in the epidemiological experiments. Naturally the opposite is the case in any experiment where the behavioral adjustments are needed for a policy to work—as in the case of increasing condom pleasure where the hope is that more people would use them. The same logic would be present for policies aimed to work through the marriage market, as laid out in Greenwood et al. (2017b).

One potential concern about this exercise is the substantial increase in the number of singles, which might be at odds with intuition. People get married for many reasons, including the desire to have children, and increased condom pleasure might not change the determination for marriage for those whom having children is very important.⁴⁵ And, fertility in Malawi is very high by Western standards. Modeling fertility explicitly goes beyond the scope of this paper, but an initial step in this direction might be accomplished by adding a value f for fertility in marriage that differs by type. Heterogeneity in this fertility benefit implies that those with strong preferences for children in marriage will not be enticed out of the marriage market, potentially reducing the responsiveness of marriage with respect to shifts in condom utility. This means that the per-period utility of marriage is now $u + l + f$. Otherwise the model is unchanged.

As a first attempt to gauge the promise of this channel, consider the following parametrization: Leave the number of types unchanged from the basic model to avoid computational complexity, but assume that those types with a high discount factor also have a higher value from fertility (i.e., have a higher f). Under

⁴⁵A referee is thanked for bringing up this point.

Table 8: CONDOMS AND FERTILITY

	<i>Benchmark</i>	<i>Better condoms</i>	<i>Benchmark (w/ fertility)</i>	<i>Better condoms (w/ fertility)</i>
p	1.4	5.5	1.4	5.5
HIV prevalence, %	10.3	15.8	10.0	13.6
Casual sex, % (of all)	15.7	34	18.5	26.9
Casual sex with condom, %	33.0	59	28.9	53.7
Singles who have casual sex, %	54.0	66	65.5	75.4
Single men, %	50	64	49.4	53.4
Single women, %	46	60	44.8	48.2

the assumption that f has a quadratic form in β , the three parameters characterizing its distribution are chosen to keep the basic calibration targets roughly unchanged relative to the benchmark model. See Appendix F.2 for more details.

Table 8 shows how close this extended version is to the original benchmark. It also shows how this new economy reacts to increased pleasure from condoms. In line with intuition, when people marry partly to have children, their marriage decision is a lot less responsive to the introduction of more pleasurable condoms. Yet, interestingly, improved condoms continue to increase the HIV rate. This is driven by singles substantially increasing their sexual activity. Overall though, as expected, the increase in the HIV rate is smaller (3.6 percentage points relative to 5.5 percentage points in the benchmark model). This shows incorporating fertility cum marriage in models could be important for analyzing HIV policy, though obviously a careful treatment would require substantially more changes to the benchmark framework.

5.3.2 Treating other Sexually Transmitted Diseases

Several policies involve the reduction in transmission risk. Large amounts of money are currently being poured into research aimed at developing a vaccine against HIV. While no successful vaccine exists yet, researchers remain hopeful that the development of a partial vaccine is possible.⁴⁶ Similarly, certain gels are

⁴⁶Promising results of a new vaccine that would reduce transmission by 30% were reported in the media as early as 2009, on September 25, the *Wall Street Journal* wrote "Vaccine Shows Promise in Preventing HIV Infections." However, additional data analysis found that the results

Table 9: TREATING OTHER STDs

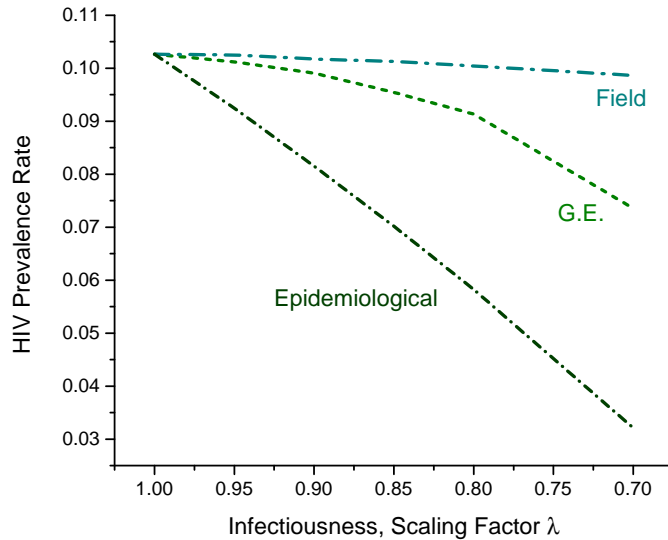
	<i>Benchmark</i>	<i>Equilibrium</i>	<i>Epidem.</i>	<i>Small Field</i>
Scaling factor	1.00	0.85	0.85	0.85
HIV prevalence, %	10.3	9.5	7.0	10.1
—Males	8.6	7.9	5.9	8.5
—Females	12.1	11.3	8.2	11.9
Casual sex, % (of all)	15.7	18.7	—	16.7
Casual sex with condom, %	33.0	27.6	—	31.2
Singles who have casual sex, %	54.0	60	—	56
Single men, %	50	52	—	51
Single women, %	46	47	—	46
Price—protected	-6.5	10	—	—
Price—unprotected	279	286	—	—
Price—long term	125	127	—	—

thought to reduce the transmission risk. Another policy that has recently been advocated is the treatment of other sexually transmitted diseases (STDs). The idea is that the presence of other STDs makes a person more susceptible to contracting HIV. Thus, treating other STDs will decrease the transmission risk, both for men and women. For example, Grosskurth et al. (1995) finds that improved STD treatment reduced HIV incidence by about 40% in rural Tanzania. Oster (2005) compares data from African countries to the United States and reaches the conclusion that treating other STDs would be an effective policy. This conclusion is based on an epidemiological simulation, which shows that differences in transmission rates (due to the existence of other STDs) can explain much of the difference in HIV rates between the United States and Sub-Saharan Africa.

Table 9 shows the simulation results for this policy in the quantitative model—see the second column. For this policy, the probability of infection is multiplied by a scaling factor $\lambda \in [0, 1]$, such that the new infection rate is $\lambda(1 - \gamma)$, for all types of sex, genders and circumcision types. As the transmission risk for both men and women declines by 15% ($\lambda = 0.85$), HIV prevalence decreases by almost one percentage point from 10.3 to 9.5%. This decrease in HIV prevalence masks the fact that, when faced with better odds when having sex, people engage in riskier behavior. The fraction of sex that is casual increases from 15.7 to 18.7%.

were statistically insignificant.

Figure 6: Treating other STDs



This is because there are more singles now, and singles have more sex. Moreover, out of the singles having sex, condom usage falls from 33 to 28%. Despite the increase in risky behavior, the policy works in the sense that HIV does fall.

Yet the behavioral changes have non-trivial effects, which can be seen as follows. Compare the results from the equilibrium experiment with the epidemiological version of the experiment in the third column of Table 9—see also Figure 6. In the epidemiological experiment, the decline in HIV prevalence is much larger, to 7%. The reason for this difference is exactly the lack of behavioral changes described above. Thus, simulations based on epidemiological experiments may significantly overstate the beneficial effects of STD treatment. This casts some doubt on Oster (2005)'s finding that STD treatment alone is able to explain much of the U.S.-S.S.A. difference in HIV as the study was based on the assumption of constant sexual behavior across countries. Allowing sexual behavior to adjust would shrink the predicted HIV gap between the United States and Sub-Saharan Africa and hence, while still leaving a sizable role for STDs, would no longer be able to explain the entire gap.⁴⁷

⁴⁷Oster (2005) argues that sexual behavior in the data is remarkably similar in the United States and Sub-Saharan Africa. However, Table II seems to indicate that behavior is somewhat more

The field experiment goes in the opposite direction: it predicts only a very small change in HIV prevalence compared to the benchmark. The reason is that, in the field experiment, the reduced number of infections does not lead to an overall decrease in the *population* prevalence rate. Therefore, it does not feed back into lower infection rates for the treated population, something that is naturally part of the full general equilibrium model.

The lack of a substantial HIV decline in the field experiment might actually be quite important. Note that eight of the nine studies of STD treatment for HIV prevention surveyed by Padian et al. (2010) delivered flat results. This seems quite puzzling, given that the theoretical effects of STD treatment are uncontroversial in the medical community. The simulations presented here highlight a novel reason for the lack of finding a large effect, namely the missing equilibrium effects in randomized field experiments. Thus, treating STDs might actually be a promising policy measure, even though it is tough to detect positive results in field experiments.

Taking stock, treating other STDs seems to decrease the overall HIV prevalence rate, even though people engage in riskier behavior in response to the lower likelihood of getting infected. This shift in behavior makes the use of a choice-theoretic general equilibrium model like the one proposed here essential. The differences arising from these behavioral responses seem to be quantitatively important.

6 The Diffusion of Better Information

HIV prevalence has fallen continuously over the last two decades—from a peak of almost 17% in 1999 to less than 10% in 2016. Section 5 suggested that the introduction of ART was not the dominant driver of the falling HIV rate in Malawi. So what explains the success story? From the policy experiments, male circumcision seems to work, yet there was no substantial change in circumcision rates in Malawi until 2010. Both the promotion of condoms and treatment of other STDs

risky in the United States, which would be in line with the model, where people in the country with lower transmission risk engage in more risky behavior.

have been part of the government's plan to fight HIV (Government of Malawi 2003), but it is unclear how widely such policies were implemented. Moreover, the policy experiments show that condom policy may backfire and the treatment of other STDs may at most have a modest effect. Thus, both policies seem unlikely to be main causes behind the HIV decline. This leaves a bit of a puzzle.

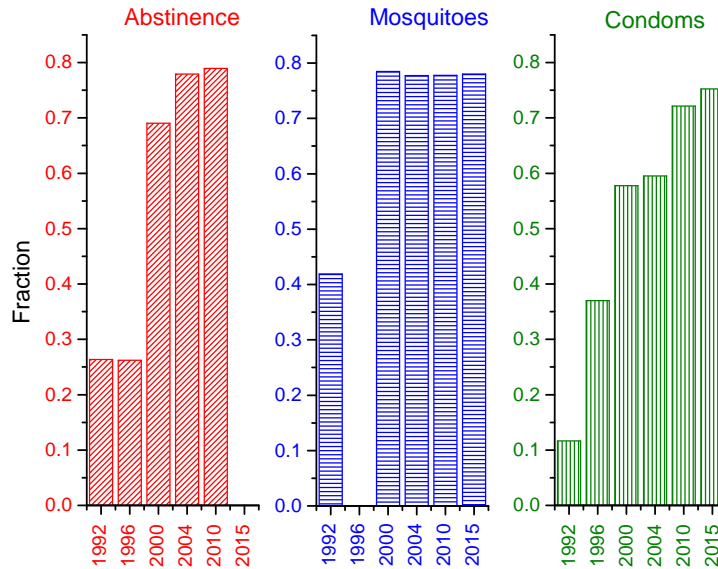
To what extent could better information have caused the decline in HIV?⁴⁸ This question is explored now. Clearly, if people are not aware of how HIV is transmitted and how they can protect themselves, they will engage in more risky activities. As information improves, behavior becomes less risky and this lowers the HIV prevalence rate. Indeed, evidence from the DHS shows that awareness of the HIV transmission mechanisms massively increased between 1992 and 2004 and has roughly stabilized since, see Figure 7.⁴⁹ Specifically, in 1992 only a quarter of the population identified correctly that abstinence reduces HIV risk, while by 2004 more than three quarters of the population did so. Similarly, in 1992 only about 10% knew that condom use reduces risk, while by 2004, 60% did. In 1992 approximately 60% erroneously believed that HIV can be transmitted through mosquito bites, which fell to about 20% by 2004. It is unclear to what extent improved awareness can be attributed to official information campaigns or whether it is the natural diffusion of information following the discovery of a new disease.⁵⁰ Without taking a stand on the underlying reason for the better information, this section asks whether the improvement in information about transmission mechanisms can explain the drop in HIV in recent years in Malawi. This is done by allowing a fraction of the population in the model to be uninformed about HIV.

⁴⁸This section is motivated by a suggestion from a referee.

⁴⁹Figure 7 is based on 6 waves of the DHS. It is similar to Figure 3.3 in Fedor (2014), but includes additional information for 2015. The two empty bars are due to the following: In 1996 the question on mosquitoes was not asked in a consistent way and in 2015 no information related to abstinence was collected. The DHS lacks data on the overall perception of the HIV transmission probability within a sexual encounter. Delavande and Kohler (2016) exploit probabilistic belief elicitation using the 2006 wave of the Malawi Longitudinal Study of Families and Health and show that individuals today systematically overestimate this probability. Lack of a longer time series precludes the investigation whether this perception changed in the late 1990s in a way that could further help explain the reversal in incidence and prevalence at that time.

⁵⁰The Malawian government disseminated a multisectoral plan called the "2000-2004 National HIV/AIDS strategic framework" in 2000. See Fedor (2014) for further details on changes in HIV knowledge in Malawi and government policies that may have contributed to the increased awareness.

Figure 7: HIV Awareness in Malawi, 1992-2015, DHS Data



To model the presence of uninformed individuals, introduce a new variable $i \in \{0, 1\}$ in the permanent-type vector, x , characterizing each individual, with the convention that $i = 1$ stands for an informed individual and $i = 0$ for an uninformed one. Informed individuals know the true structure of transmission. Uninformed people do not. The survey evidence indicates the presence of individuals who are not aware that either condoms or abstinence can protect them from HIV infection. Therefore, assume that uninformed singles believe that the odds of contracting HIV are the same as in the short-term unprotected market, regardless of their sexual behavior. That is, individuals who either enter the short-term market for protected sex or who remain abstinent think that they are just as likely to catch HIV as someone in the short-term unprotected market. In the absence of any other anchor for their beliefs about unprotected sex, for simplicity assume that they assign correct beliefs to unprotected sex. Note that all sex in the long-term market is unprotected so nothing changes here. Finally, even though the survey did not directly question individuals about circumcision, it is unlikely that uninformed individuals were aware of its protective effects. So assume that uninformed individuals do not assign protective power to circumcision.

The adjustments to the model required to incorporate the presence of uninformed

people are relatively modest. The value functions for uninformed individuals are similar to those for informed ones, with just a few changes. These adjustments imply different behavior for them. Specifically, for a healthy uninformed old individual the value of abstinence and short-term sex in equations (1) and (3) are adjusted by replacing the infection probabilities by those in the short-term unprotected market and by setting $\chi(c) = 1$, for all c , so that individuals do not take account of circumcision. The adjustment to the analogs for young healthy individuals are identical. For uninformed old individuals in the long-term market, (7) is adjusted so that they do not take into account the protective effects of circumcision for themselves or for their circumcised partner. Again, the same is done for young individuals in the long-term market. Given the decision rules for uninformed people, the determination of the equilibrium incidence and prevalence rates is governed by the correct transmission parameters. The exact equations are presented in Appendix F.1.

The model with uninformed people can now be used to explore to what extent changes in information may have been responsible for the decline in HIV over time. Focus on two years: 1996 and (as before) 2004. Since awareness was relatively high and stable from 2004 onward, the full-information model is a good approximation of reality for that year.⁵¹ For 2004, it will be assumed that 60 percent of the population is uninformed in line with the data.⁵² Table 10 gives the results and compares them to the data for the same two years. The data moments are from two waves of the DHS (1996 and 2004) except for the HIV incidence and prevalence numbers.⁵³

The table shows that adding lack of awareness about the HIV transmission mechanism indeed raises the overall prevalence and incidence rates quite a bit. The

⁵¹Of course, as Figure 7 shows, even in 2004 there are still some uninformed people. The benchmark model was deliberately kept the same as in the previous working paper version (Greenwood et al. 2017a) as this gives an additional opportunity to examine the performance of the model with non-targeted moments.

⁵²As Figure 7 shows, in 1996, 63% of the population did not know that a condom is effective, 74% did not know that abstinence helps, and (in 1992) 58% erroneously believed that HIV can be transmitted through mosquitoes.

⁵³The HIV incidence and prevalence numbers are from UNAIDS, see also Footnote 35. This departure from Table 2 is due to the fact that the 1996 DHS did not collect biomarker information and hence no HIV prevalence rate can be calculated. Further, none of the DHS has incidence information.

Table 10: IMPROVED INFORMATION—MODEL VS. DATA

	Data			Model		
	1996	2004	Δ in p.p.	1996	2004	Δ in p.p.
HIV prevalence, %	14.6	13.5	-1.1	14.4	10.3	-4.1
HIV incidence, per 1,000 healthy	20.4	11.1	-9.3	17.3	11.7	-5.6
Sex that is casual, %	24	18	-6	20	16	-4
Condom use in casual sex, %	29	39	10	28	33	5
Fraction of singles, %	31	33	2	47	48	1
Singles who had casual sex, %	60	47	-13	67	54	-13

implied decline in the HIV rate is in fact larger than the one observed in the data. This is not surprising since the model results are based on a steady-state comparison, while in reality the transitional dynamics would likely take some time to work out. On this, note that the prevalence rate simply cannot fall very quickly in the real world because it is a stock variable. It takes time for the reduction in new infections (incidence) to cumulate through the economy. A more meaningful comparison are thus the incidence rates. Incidence rates actually fell very quickly over this time period from 20.4 new infections annually per 1,000 healthy people in 1996 to 11.1 by 2004. The model implies that a reduction in the fraction of the population that is uninformed from 60% to zero will decrease incidence rates by a similar order of magnitude from 17.3 to 11.7 per 1,000 healthy people. Therefore improved information may well have played a key role in reducing HIV in Malawi.

The fall in incidence rates is due to safer sexual choices. The adjustments in terms of sexual behavior are also remarkably close to the data. In the model, the fraction of casual sex falls by 4 percentage points (p.p.), condom use increases by 5 p.p., the fraction of singles declines by 1 p.p., and among them 13 p.p. fewer have casual sex. The data shows a very similar trend towards less risky behavior. From 1996 to 2004 the fraction of casual sex fell by 6 p.p., condom use increased by 10 p.p., the fraction of singles fell by 2 p.p., and among them 13 p.p. fewer have casual sex.

Finally, given that in reality even in 2004 not all people were perfectly informed, one may ask the question to what extent the assumption of perfect awareness for the benchmark model biases the policy experiments in Section 5. Put differently,

one may want to know what the effect of policies would have been in a world with more uninformed people. This question is addressed by implementing the same policy experiments for the 1996 model (i.e., the one with 60% uninformed people). This is reported in Table 13 in Appendix F.1. The table shows that results are qualitatively similar to those in the full information model. At the same time, all policies are quantitatively somewhat more effective in the model with uninformed people. For example circumcising all men would lower the HIV rate by 68% relative to only 46% in the baseline model. In the case of ART the positive effect is seen earlier; i.e., now treating a third of the population ($q = 0.01$) leads to a decline in the HIV rate. The higher effectiveness of the policies is not that surprising, though. Recall that in the model with perfect information people dampen the effectiveness of the policies by responding with riskier behavior. Uninformed people do not react as much because they do not understand the true trade-offs between different sexual behaviors. Thus, the higher the fraction of uninformed people, the better all policies work in reducing HIV.

7 Conclusions

In Malawi about 12 percent of the population has the HIV/AIDS virus. Roughly 18 percent of sex is casual and a condom is used a third of the time. An equilibrium search model is constructed to analyze the Malawian HIV/AIDS epidemic. At the heart of the model is homo economicus. Specifically, it is presumed that economic man or woman searches for the type of sexual activity that they desire to engage in, while rationally taking into consideration the risks of this activity. Some people will select stable long-term relationships, others may choose more fleeting ones. Condoms may or may not be used in these more ephemeral encounters, depending on the participants' mutual desires. The number of such encounters is partially under people's control. All of these choices crucially affect the spread of HIV/AIDS in society.

The aim of the analysis is to provide a toolbox that allows the study of various interventions, identifies where behavioral change and equilibrium effects might be important, and thereby pinpoints areas where further and deeper exploration

might be most warranted.

The theoretical model developed is calibrated to capture some of the salient features of the Malawian HIV/AIDS epidemic. The framework can match both targeted and non-targeted statistics regarding sexual behavior and HIV/AIDS in Malawi as well as some cross-country data. The benchmark simulation is then used to undertake some policy interventions. The quantitative results suggest that policy analysis of HIV/AIDS interventions is complicated. In particular, some policies may backfire and actually increase HIV. The simulations also rationalize some puzzling results in previous works.

Specifically, it is shown that ART was unlikely the main driver behind the decline in HIV over time. ART can work well, but only if a large fraction of the population is treated. Second, male circumcision seems a useful policy with effects that are likely much larger than one would expect based on field experiments alone. Third, better condoms could potentially backfire and lead to more HIV. Fourth, the beneficial effects of treating other STDs are much smaller than is sometimes argued based on epidemiological studies. At the same time, this study resolves a puzzle surrounding the lack of positive effects in previous field experiments due to the difficulty of capturing the salubrious spillover effects on partners that operate to reduce the equilibrium HIV prevalence rate. Finally, the analysis also explores the likelihood that the populace was not fully knowledgeable in the early stages of the epidemic about how HIV is transmitted. As a result, people who were uninformed did not move into safer sexual activity. It is shown that as information spreads about the HIV transmission mechanism the rate of HIV falls. Thus, the diffusion of better information was probably the most important driver of the HIV decline in Malawi over time.

The framework developed here contains several dimensions of heterogeneity. One can, however, think of different margins that were not explored in the current paper. For example, HIV prevalence rates in urban areas tend to be higher than in rural regions. This may have to do with different search costs in the two localities.⁵⁴ On this, the impact of migration between the rural and urban areas may be important to consider. Further, the theoretical model does not allow

⁵⁴Indeed, running a counterfactual in which search costs in the model are increased leads to a lower HIV prevalence rate.

for concurrent relationships, while in reality extra-marital affairs and polygyny may matter for the spread of the virus. Such questions can be tackled in future research.

References

- Ahuja, Amrita, Brian Wendell, and Eric Werker. 2009. "Male circumcision and AIDS: the macroeconomic impact of a health crisis." *Harvard Business School Working Paper*, vol. 07-025.
- Allen, Susan, Jareen Meinzen-Derr, Michele Kautzman, Isaac Zulu, Stanley Trask, Ulgen Fideli, Rosemary Musonda, Francis Kasolo, Feng Gao, and Alan Hanworth. 2003. "Sexual Behaviour of HIV Discordant Couples after HIV Counseling and Testing." *AIDS* 17 (5):733–740.
- Anderson, Roy M., and Robert M. May. 1992. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, Oxford.
- Ashraf, Nava, Oriana Bandiera, and Kelsey Jack. 2014. "No Margin, no Mission? A Field Experiment on Incentives for Public Service Delivery." Unpublished Manuscript, London School of Economics.
- Auld, M. Christopher. 2006. "Estimating Behavioral Response to the AIDS Epidemic." *B.E. Journals in Economic Analysis and Policy: Contributions in Economic Analysis & Policy* 5 (1):12.
- Auvert, Bertran, Dirk Taljaard, Emmanuel Lagarde, Jolle Sobngwi-Tembekou, Rmi Sitta, and Adrian Puren. 2005. "Randomized, Controlled Intervention Trial of Male Circumcision for Reduction of HIV Infection Risk: The ANRS 1265 Trial." *PLoS Med* 2 (11):1112–1122.
- Baeten, Jared M., Barbra A. Richardson, Ludo Lavreys, Joel P. Rakwar, Kishorchandra Mandaliya, Job J. Bwayo, and Joan K. Kreiss. 2005. "Female-to-Male Infectivity of HIV-1 among Circumcised and Uncircumcised Kenyan Men." *Journal of Infectious Diseases* 191 (4):546–553.
- Bailey, R. and Neema, S., and Richard Othieno. 1999. "Sexual Behavior and Other HIV Risk Factors in Circumcised and Uncircumcised Men in Uganda." *Journal of Acquired Immune Deficiency Syndrome* 22:294–301.

- Bailey, R.C., S. Moses, C.B. Parker, K. Agot, I. Maclean, J.N. Krieger, C.F.M. Williams, R.T. Campbell, and J.O. Ndinya-Achola. 2007. "Male circumcision for HIV prevention in young men in Kisumu, Kenya: a randomised controlled trial." *The Lancet* 369 (9562): 643–656.
- Beksinska, Mags E., Jennifer A. Smit, and Joanne E. Mantell. 2012. "Progress and challenges to male and female condom use in South Africa." *Sexual Health* 9 (1): 51–58.
- Bowie, Cameron. 2006. "The Burden of Disease in Malawi." In *The Epidemiology of Malawi*, edited by Eveline Geubbels and Cameron Bowie. College of Medicine, University of Malawi.
- Bracher, Michael, Gigi Santow, and Susan Cotts Watkins. 2004. "Assessing the Potential of Condom Use to Prevent the Spread of HIV: A Microsimulation Study." *Studies in Family Planning* 35 (1):48–64.
- Canning, David. 2006. "The Economics of HIV / AIDS in Low-Income Countries: The Case for Prevention." *Journal of Economic Perspectives* 20 (3):121–142.
- Chimbiri, Agnes M. 2007. "The Condom is an 'Intruder' in Marriage. Evidence from Rural Malawi." *Social Science and Medicine* 64 (5):1102–1115.
- Clark, Samuel J., and Jeff Eaton. 2008. "Demographic Consequences of HIV Epidemics and Effects of Different Male Circumcision Intervention Designs: Suggestive Findings from Microsimulation." Working Paper no. 85, Center for Statistics and the Social Sciences, University of Washington.
- Cohen, Craig R., Michele Montandon, Adam W. Carrico, Stephen Shiboski, Alan Bostrom, Alfredo Obure, Zachary Kwena, Robert Bailey, Rosemary Nguti, and Elizabeth A. Bukusi. 2009. "Association of Attitudes and Beliefs towards Antiretroviral Therapy with HIV-Seroprevalence in the General Population of Kisumu, Kenya." *PLOS One*, vol. 4.
- Cohen, M.S. et al. 2011. "Prevention of HIV-1 Infection with Early Antiretroviral Therapy." *The New England Journal of Medicine* 365 (6): 493–505.
- Crepaz, Nicole, Trevor A. Hart, and Gary Marks. 2004. "Highly Active Antiretroviral Therapy and Sexual Risk Behavior." *Journal of the American Medical Association* 292 (2): 224–236.

- Delavande, Adeline, and Hans-Peter Kohler. 2009. "Subjective Expectations in the Context of HIV/AIDS in Malawi." *Demographic Research* 20:817–874.
- . 2016. "HIV/AIDS-related Expectations and Risky Behavior in Malawi." *Review of Economic Studies* 83:118–164.
- De Walque, D., W.H. Dow, R. Nathan, and et al. 2012. "Incentivizing safe sex: A randomized trial of conditional cash transfers for HIV and sexually transmitted infection prevention in rural Tanzania." *BMJ Open* 2, no. e000747.
- De Walque, Damien. 2007. "How Does the Impact of an HIV/AIDS Information Campaign Vary with Educational Attainment? Evidence from Rural Uganda." *Journal of Development Economics* 84 (2):686–714.
- . 2012. "Economics of HIV/AIDS: transmission, treatment, and prevention." In *Rethink HIV. Smarter Ways to Invest in Ending HIV in Sub-Saharan Africa*, edited by Bjorn Lomborg. Cambridge University Press.
- De Walque, Damien, Harounan Kazianga, and Mead Over. 2012. "Antiretroviral Therapy Perceived Efficacy and Risky Sexual Behaviors: Evidence from Mozambique." *Economic Development and Cultural Change* 61 (1): 97–126.
- Dupas, P. 2011. "Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya." *American Economic Journal: Applied Economics* 3 (1):1–34.
- Eckhout, Jan, and Philipp Kircher. 2010. "Sorting and Decentralized Price Competition." *Econometrica* 78 (2):539–574.
- Fedor, Theresa. 2014. "Changes in HIV/AIDS Knowledge, Attitudes and Behavior in Malawi." Dissertation, University of Pennsylvania.
- Fenichel, Eli P., Carlos Castillo-Chavez, M. G. Ceddiac, Gerardo Chowell, Paula A. Gonzalez Parra, Graham J. Hickling, Garth Hollowayc, Richard Horan, Benjamin Morin, Charles Perrings, Michael Springborn, and Ve. 2011. "Adaptive human behavior in epidemiological models." *Proceedings of the National Academy of Science (PNAS)* 108 (15):6306–6311.
- Friedman, Willa. 2015, January. "Antiretroviral Drug Access and Behavior Change." Unpublished Manuscript, University of Houston.

- Gautier, Pieter, Michael Svarer, and Coen N. Teulings. 2010. "Marriage and the City: Search Frictions and Sorting of Singles." *Journal of Urban Economics* 67 (2):206–218.
- Gertler, Paul, Manisha Shah, and Stefano M. Bertozzi. 2005. "Risky Business: The Market for Unprotected Commercial Sex." *Journal of Political Economy* 113 (3):518–550.
- Government of Malawi. 2003. "National HIV/AIDS Policy – A Call for Renewed Action." Technical Report, National AIDS Commission, Office of the President and Cabinet.
- Gray, Ronald H., Godfrey Kigozi, David Serwadda, Frederick Makumbi, Stephen Watya, Fred Nalugoda, and Noah Kiwanuka. 2007. "Male Circumcision for HIV Prevention in Men in Rakai, Uganda: A Randomized Trial." *Lancet* 369 (9562):657–666.
- Gray, Ronald H., Maria J. Wawer, Ron Brookmeyer, Nelson K. Sewankambo, David Serwadda, Fred Wabwire-Mangen, Tom Lutalo, Xianbin Li, Thomas vanCott, Thomas C. Quinn, and the Rakai Project Team. 2001. "Probability of HIV-1 Transmission per Coital Act in Monogamous Heterosexual, HIV-1-Discordant Couples in Rakai, Uganda." *Lancet* 357 (9263):1149–1153.
- Greenwood, Jeremy, Philipp Kircher, Cezar Santos, and Michèle Tertilt. 2013. "An Equilibrium Model of the African HIV/AIDS Epidemic." Working paper 18953, National Bureau of Economic Research.
- . 2017a. "An Equilibrium Model of the African HIV/AIDS Epidemic." Working paper 601, Rochester Center for Economic Research.
- Greenwood, Jeremy, Philipp Kircher, Cezar Santos, and Michele Tertilt. 2017b. "The Role of Marriage in Fighting HIV: A Quantitative Illustration for Malawi." *American Economic Review P&P* 107 (5): 158–162.
- Grosskurth, H., F. Mosha, J. Todd, E. Mwijarubi, A. Klokke, K. Senkoro, P. Mayaud, J. Changalucha, A. Nicoll, J. ka Gina, G. and Newell, K. Mugeye, D. Mabey, and R. Hayes. 1995. "Impact of Improved Treatment of Sexually Transmitted Diseases on HIV Infection in Rural Tanzania: Randomized Controlled Trial." *Lancet* 346 (8974):530–536.

- Group, The TEMPRANO ANRS 12136 Study. 2015. "A Trial of Early Antiretrovirals and Isoniazid Preventive Therapy in Africa." *New England Journal of Medicine* 373 (9): 808–822. PMID: 26193126.
- Hethcote, Herbert W. 2000. "The Mathematics of Infectious Diseases." *SIAM Review* 42 (4):599–653.
- Iorio, Daniela, and Raul Santaaulalia-Llopis. 2011. "Education, HIV Status, and Risky Sexual Behavior: How Much Does the Stage of the HIV Epidemic Matter?" Unpublished Manuscript, Washington University.
- Jacquet, Nicolas L., and Serene Tan. 2007. "On the Segmentation of Markets." *Journal of Political Economy* 115 (4):639–664.
- Johnson, Leigh Francis. 2008. "The interaction between HIV and other sexually transmitted infections in South Africa: A model-based evaluation." Ph.D. diss., Department of Actuarial Science, University of Cape Town.
- Jovanovic, Boyan. 1979. "Job Matching and the Theory of Turnover." *Journal of Political Economy* 87 (5):972–990.
- Kim, Hyuncheol Bryant, booyuel Kim, Cristian Pop-Eleches, and Jaehyun Jung. 2018, February. "Male Circumcision, Peer Effects and Risk Compensation." Unpublished Manuscript, Department of Policy Analysis and Management, Cornell University.
- Klein, Eili, Ramanan Laxminarayan, David L. Smith, and Christopher A. Gilligan. 2007. "Economic incentives and mathematical models of disease." *Environment and Development Economics* 12 (5):707–732.
- Kremer, Michael. 1996. "Integrating Behavioral Choice into Epidemiological Models of AIDS." *The Quarterly Journal of Economics* 111 (2): 549–573 (May).
- Kremer, Michael, and Charles Morcom. 1998. "The Effect of Changing Sexual Activity on HIV Prevalence." *Mathematical Biosciences* 151 (1):99–122.
- Lakdawalla, Darius, Neeraj Sood, and Dana Goldman. 2006. "HIV Breakthroughs and Risky Sexual Behavior." *Quarterly Journal of Economics* 121 (3):1063–1102.
- Low-Beer, D., and R.L. Stoneburner. 1997. "An age- and sex-structured HIV epi-

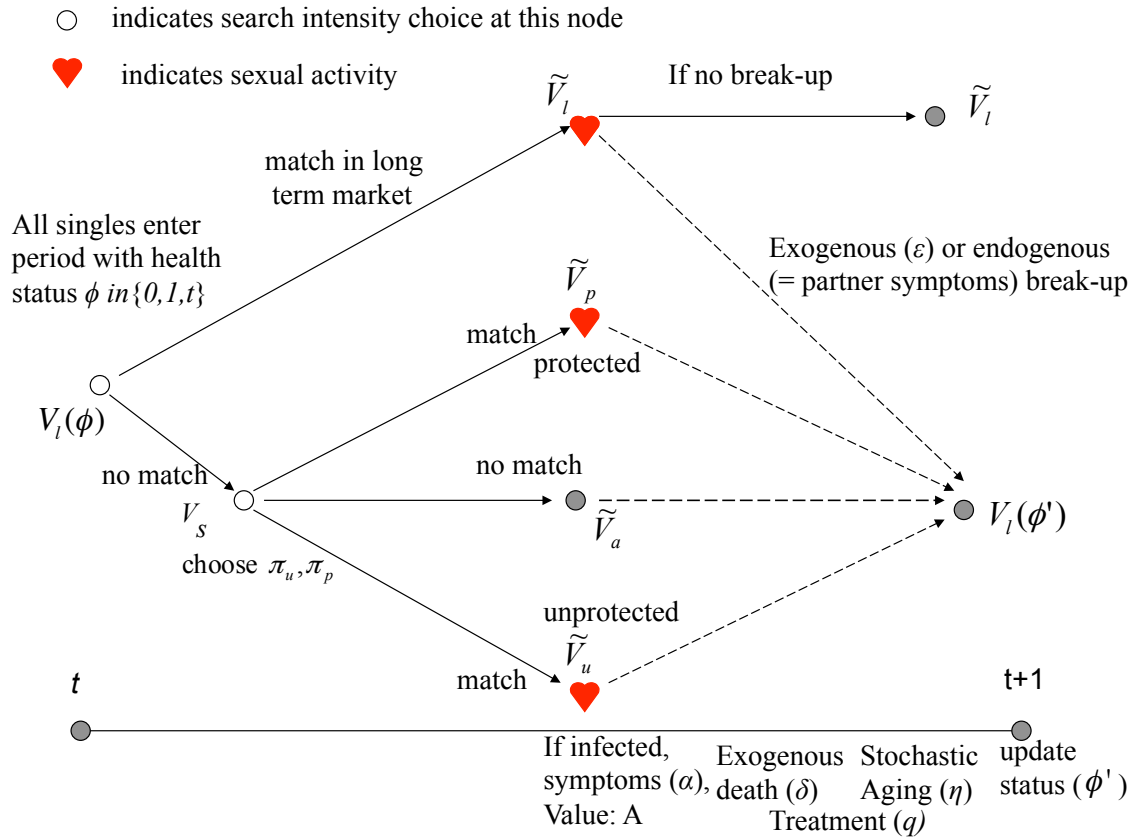
- demiological model: features and applications." *Bulletin of the World Health Organization* 75 (3): 213–221.
- Lowrance, David. 2007. "Lower Early Mortality Rates Among Patients Receiving Antiretroviral Treatment at Clinics Offering Cotrimoxazole Prophylaxis in Malawi." *JAIDS Journal of Acquired Immune Deficiency Syndromes* 46:56–61.
- Magalhães, Leandro De, and Raül Santaaulàlia-Llopis. 2018. "The consumption, income, and wealth of the poorest: An empirical analysis of economic inequality in rural and urban Sub-Saharan Africa for macroeconomists." *Journal of Development Economics* 134:350–371.
- Magruder, J. 2011. "Marital Shopping and Epidemic AIDS." *Demography* 48 (4):1401–1428.
- Nicolosi, Alfredo, Maria Lea Correa Leite, Massimo Musicco, Claudio Arici, Giovanna Gavazzeni, and Adriano Lazzarin. 1994. "The Efficiency of Male-to-Female and Female-to-Male Sexual Transmission of the Human Immunodeficiency Virus: A Study of 730 Stable Couples." *Epidemiology* 5 (6):570–575.
- Nzyuko, S., P. Lurie, W. McFarland, W. Leyden, D. Nyamwaya, and J.S. Mandel. 1997. "Adolescent sexual behavior along the Trans-Africa Highway in Kenya." *AIDS Suppl* 1:S21–26.
- Oster, Emily. 2005. "Sexually Transmitted Infections, Sexual Behavior, and the HIV/AIDS Epidemic." *Quarterly Journal of Economics* 120 (2):467–515.
- . 2012. "Routes of Infection: Exports and HIV Incidence in Sub-Saharan Africa." *Journal of the European Economic Association* 10 (5): 1025–1058.
- Padian, Nancy S., S. I. McCoy, J. E. Balkus, and J. N. Wasserheit. 2010. "Weighing the Gold in the Gold Standard: Challenges in HIV Prevention Research." *AIDS* 24 (5):621–635.
- Padian, Nancy S., Stephen C. Shiboski, and Nicholas P. Jewell. 1991. "Female-to-Male Transmission of Human Immunodeficiency Virus." *Journal of the American Medical Association* 266 (12):1664–1667.
- Philipson, Thomas J., and Richard A. Posner. 1993. *Private Choices and Public Health - The AIDS Epidemic in an Economic Perspective*. Harvard University Press, Cambridge, Massachusetts.

- Porco, T., J.N. Martin, K.A. Page-Shafer, A. Cheng, E. Charlebois, R.M. Grant, and D.H. Osmond. 2004. "Decline in HIV infectivity following the introduction of highly active antiretroviral therapy." *AIDS* 18 (1): 81–88.
- Poulin, Michelle. 2007. "Sex, Money, and Premarital Partnerships in Southern Malawi." *Social Science and Medicine* 65 (11):2383–2393.
- Ray, Maile et al. 2010. "The effect of combined antiretroviral therapy on the overall mortality of HIV-infected individuals." *AIDS* 24 (1): 123–137.
- Reniers, Georges. 2003. "Divorce and Remarriage in Rural Malawi." *Demographic Research* 1 (6):175–206.
- Rolison, Jonathan J., Yaniv Hanoch, Stacey Wood, and Pi-Ju Liu. 2013. "Risk-Taking Differences Across the Adult Life Span: A Question of Age and Domain." *Journals of Gerontology, Series B: Psychological Sciences and Social Sciences* 69 (6): 870–880 (October).
- Santaaulalia-Llopis, Raul. 2008. "Aggregate Effects of AIDS on Development." Washington University in St. Louis, Unpublished Manuscript.
- Seed, J., S. Allen, T. Mertens, E. Hudes, A. Serufulira, M. Carael, E. Karita, P. Van de Perre, and F. Nsengumuremyi. 1995. "Male Circumcision, Sexually Transmitted Disease, and Risk of HIV." *Journal of Acquired Immune Deficiency Syndrome*, pp. 883–890.
- Swidler, Ann, and Susan Cotts Watkins. 2007. "Ties of Dependence: AIDS and Transactional Sex in Rural Malawi." *Studies in Family Planning* 38 (3):147–162.
- Thornton, Rebecca. 2008. "The Demand for and Impact of Learning HIV Status." *American Economic Review* 98 (5):1829–1863.
- Toxvaerd, Flavio. 2010. "Recurrent Infection and Externalities in Prevention." CEPR Discussion Paper 8112.
- UNAIDS. 2016. "Prevention Gap Report." Technical Report, UNAIDS, Geneva, Switzerland.
- UNAIDS Joint United Programme on HIV/AIDS, ed. 2007. *Modelling the Expected Short-Term Distribution of Incidence of HIV Infections by Expose Group - Manual*.

- Undie, Chi-Chi, Joanna Crichton, and Eliya Zulu. 2007. "Metaphors We Love By: Conceptualizations of Sex among Young People in Malawi." *African Journal of Reproductive Health* 11 (3):221–235.
- U.S. Census Bureau, ed. 2004. *The AIDS Pandemic in the 21st Century. International Population Reports WP/02-2*. U.S. Government Printing Office, Washington DC.
- Wawer, Maria J., Ronald H. Gray, Nelson K. Sewankambo, David Serwadda, Xianbin Li, Oliver Laeyendecker, Noah Kiwanuka, Godfrey Kigozi, Mohammed Kiddugavu, Thomas Lutalo, Fred Nalugoda, and Fred a Wabwire Mangen. 2005. "Rates of HIV-1 Transmission per Coital Act, by Stage of HIV-1 Infection, in Rakai, Uganda." *Journal of Infectious Diseases* 191 (9):1403–1409.
- Weller, Susan C. 1993. "A Meta-Analysis of Condom Effectiveness in Reducing Sexually Transmitted HIV." *Social Science and Medicine* 36 (12):1635–1644.
- Wellings, Kaye, Julia Field, Anne M. Johnson, Jane Wadsworth, and Sally Bredshaw, eds. 1994. *Sexual Behavior in Britain. The National Survey of Sexual Attitudes and Lifestyles*. Penguin Books Ltd, London, England.
- Williams, B.G, J.O. Lloyd-Smith, Eleanor Gouws, Catherine Hankins, W.M. Getz, J. Hargrove, I. de Zoysa, Christopher Dye, and B. Avert. 2006. "The Potential Impact of Male Circumcision on HIV in Sub-Saharan Africa." *PLoS Medicine* 3 (7): 1032–1040.
- Wilson, David P. 2012. "HIV Treatment as Prevention: Natural Experiments Highlight Limits of Antiretroviral Treatment as HIV Prevention." *PLOS Medicine* 9, no. 7.
- World Health Organisation, ed. 2008. *The Global Burden of Diseases: 2004 update*. WHO Press.
- World Health Organization. 2016. "PROGRESS REPORT 2016: PREVENT HIV, TEST AND TREAT ALL." Technical Report, World Health Organization, Geneva, Switzerland.
- Young, Alwyn. 2005. "The Gift of the Dying: The Tragedy of AIDS and the Welfare of Future African Generations." *Quarterly Journal of Economics* 120 (2): 423–66 (May).

A Appendix—Timing of Events

Figure 8: Timing of Events



B Appendix—Data

Most of the empirical moments are based on information collected from the interviews of individuals conducted for the Malawi Demographic and Health Survey (MDHS) in 2004, carried out by the Malawi National Statistical Office. In this nationally representative survey 11,698 women aged 15 to 49 and 3,261 men aged 15 to 54 were interviewed. Means are calculated using sample weights. For several figures means are calculated by age. Since men are underrepresented in the survey, separate means are calculated by sex, and then averaged. Whenever sources other than the MDHS are used, it will be indicated. More details on each figure follow. For the interested reader the details also include the variable names corresponding to each question.

- Figure 1: Fraction ever Married—Model vs. Data
The fraction of ever-married people is derived by dividing the number of people who either are currently married (including cohabitation) or have been formerly married by all people. The corresponding question is “Have you ever been married or lived with a man/woman” (MDHS 2004: v/mv502).
- Figure 9: HIV Rate—Men vs. Women, Model vs. Data
In order to calculate the HIV rates by age (MDHS 2004: v012/mv012) and gender, individual information from the MDHS 2004 is matched with the HIV test results (MDHS 2004: hiv03) for those people who agreed on doing the test along with the interview (since not everyone agreed, the sample size is smaller here: 2,404 men and 2,864 women). The resulting HIV rates are smoothed using a third-order polynomial.
- Figure 10: Sexual Behavior by Age—Model vs. Data
Singles: Those men and women who reported that they have never been married or are widowed, divorced (living or not living together) are defined as singles (MDHS 2004: v/mv501).
Casual sex: To identify the fraction of sex that occurs in casual relationships, all men and women are considered who had sex in the last year (MDHS 2004: v/mv529). These people are asked with whom they had sex (MDHS 2004: v/mv767a). They are also asked whether they had sex with a second

(MDHS 2004: v/mv761b, v/mv767b) and third (MDHS 2004: v/mv761c, v/mv767c) partner. If one of the sex partners was not the spouse or cohabiting partner, then the sex in the last year is categorized as casual sex. Men in addition are asked whether they have ever paid for sex (MDHS 2004: mv792). Those men who have paid for sex in the last year are also defined as being active in the short-term market.

- Figure 11: Deaths by HIV/AIDS by Age—Model vs. Data
The data on deaths caused by HIV/AIDS are taken from Bowie (2006), pages 31-42. He reports the fraction of HIV/AIDS related deaths by age groups, based on the WHO Global Burden of Disease Malawi from 2002.
- Table 1: Parameters Chosen Outside the Model
All sources are described in the text.
- Table 3: Targeted Moments
The data on the prevalence of HIV/AIDS in Malawi derive from the Demographic and Health Surveys' (MDHS) Final Survey for Malawi in 2004. See MDHS (2004, Table 12.3). The fraction of sex that is casual is the proportion of people—averaged across men and women—who had sex with a non-marital, non-cohabiting partner during the last year, conditional on being sexually active, and is taken from MDHS (2004, Table 11.9). Condom usage for short-term sex also derives from MDHS (2004, Table 11.9)—and is averaged across men and women. The fraction of singles who have casual sex is reported in MDHS (2004, Tables 6.71 and 6.72) and corresponds to the weighted average of never married and divorced/separated/widowed men and women. The proportion of the population that is single is contained in MDHS (2004, Table 6.1), where single is interpreted as anyone who is not currently married nor cohabiting, averaged across men and women. The fraction of males and females that has ever been married by a certain age is the same as in Figure 1. The World Health Organisation (2008) reports that 29% of all deaths in Malawi in 2004 were due to HIV/AIDS.
- Figure 2 and Table 4: The cross-country circumcision data comes from Ahuja, Wendell, and Werker (2009). The statistics for HIV rates come from UNAIDS while the numbers for GDP per capita and ART coverage come from

the World Bank Development Indicators. The rates for syphilis seropositivity relates to data among antenatal care attendees from the WHO Global Health Observatory. The fractions of populations of different religions are given by the Global Religious Futures Project of the Pew Research Center. Condom prices for different countries are reported in the Global Directory of Condom Social Marketing Projects and Organisations (UNAIDS).

- Figure 3: HIV is defined as “Prevalence of HIV, total (% of population ages 15-49).” ART is defined as “Antiretroviral therapy coverage (% of people living with HIV).” HIV data comes from UNAIDS and ART coverage is taken from the World Development Indicators.
- Table 10: Incidence and prevalence numbers are taken from UNAIDS. All numbers on sexual behavior are computed from the MDHS. The 2004 numbers are identical to those in Table 3 and were calculated as described above. The numbers for 1996 were computed in exactly the same way using data from the MDHS 1996 instead. Model numbers are based on own model simulations.
- Figure 7 is based on the MDHS for the years 1992, 1996, 2000, 2004, 2010, and 2015. The question asked is “What can a person do [to avoid getting AIDS]?” and then various options are given. For 1996 the question on mosquitoes was not asked in a consistent way and in 2015 no information related to whether abstinence reduces HIV risk was collected. Fractions were computed based on the entire sample (males and females); sample weights were used. The exact list of variables used is as follows. Abstinence: women: s808a (1992), qa509a (1996), v754b (2000-2010), men: mv754b (1992), qma509a (1996), mv754b (2000-2010). Condom use: women: s808c (1992), qa504 (1996), v754c (2000-2015) men: mv754c (1992), qma504 (1996), mv754c (2000-2015). Mosquito: women: s804g (1992), v754jp (2000-2015) men: sm504g (1992), mv754jp (2000-2015).

C Appendix—Robustness

This appendix provides some sensitivity analysis regarding the parameters estimated on Section 4. Recall that 11 parameters were chosen by fitting the model to a specific set of data moments from Malawi. These are listed on the different rows of Tables 11 and 12. Each of these two tables has three columns besides the first that lists the parameters. The column labeled “HIV—benchmark” provides the HIV prevalence rate when the parameter of each corresponding row is changed by 1% (Table 11) or 10% (Table 12).⁵⁵ The column “ Δ HIV—circumcision (50%)” reports the change in HIV rate under the intervention that circumcises 50% of the males in the economy. Finally, the last column (Δ HIV—ART ($q = 5\%$)) presents the change in HIV rate when the infected have a 5% probability of receiving ART in each period.

Table 11 shows that the benchmark is quite robust when the parameters are changed by 1%. The HIV prevalence rate is always remarkably close to the 10.3% found in the benchmark calibration. Moreover, the results from the two main policy experiments (male circumcision and ART) are also very close to the changes found in the benchmark. Juxtapose these numbers with the ones reported in Table 12, in which each parameter is changed by 10%. The percentage change now is considerably larger. Correspondingly, the HIV prevalence rate now changes compared with the main calibration. This suggests that, in order to fit the moments targeted in the calibration, the parameters should be close to the ones found in the estimation. At the same time, the percentage changes in the policy experiments are remarkably similar, even if individual parameters are changed by 10%.

⁵⁵To be precise, the rows for the discount factors (β_{max} and β_{min}) report changes on the discount rates $\rho = (1 - \beta)/\beta$.

Table 11: ROBUSTNESS—1%

	HIV—benchmark	Δ HIV—circumcision (50%)	Δ HIV—ART ($q = 5\%$)
Main calibration	10.3	-1.2	-0.8
p	10.3	-1.2	-0.7
u	10.2	-1.2	-1.1
l	10.1	-1.2	-1.0
β_{max}	10.3	-1.2	-0.8
β_{min}	10.4	-1.3	-0.8
ι_{change}	10.2	-1.2	-0.7
A	10.3	-1.2	-0.9
η	10.2	-1.2	-0.7
ω_s	10.2	-1.2	-0.9
ω_l	10.3	-1.2	-0.7
κ	10.3	-1.2	-0.8

Table 12: ROBUSTNESS—10%

	HIV—benchmark	Δ HIV—circumcision (50%)	Δ HIV—ART ($q = 5\%$)
Main calibration	10.3	-1.2	-0.8
p	10.3	-1.2	-0.8
u	9.1	-1.5	0.0
l	8.8	-1.3	0.3
β_{max}	10.3	-1.2	-0.8
β_{min}	11.4	-1.5	-1.8
ι_{change}	9.0	-0.5	0.5
A	10.3	-1.2	-0.8
η	9.8	-1.1	-0.3
ω_s	9.8	-1.3	-0.5
ω_l	10.9	-1.3	-1.3
κ	10.1	-1.2	-0.7

D Appendix—Theory

D.1 Value Functions for Young Individuals, $d = \iota$

The value functions for young individuals follow a similar structure as those for old individuals, namely equations (1) to (6). The required adjustments were outlined in the main body in connection with (1).

In particular, for young abstinent individuals of health status ϕ the analog to (1) replaces the high discount factor, β , with the low discount factor, ι , and treats continuation values as the average of the continuation with a low and a high discount factor, so that

$$\tilde{V}_a^\iota(\phi, x) = \ln(y) + \alpha_\phi \iota A + [1 - \alpha_\phi] \iota \left\{ \begin{array}{l} Q(\phi)[\eta V_t^\beta(t, x) + (1 - \eta)V_t^\iota(t, x)] \\ + [1 - Q(\phi)][\eta V_t^\beta(\phi, x) + (1 - \eta)V_t^\iota(\phi, x)] \end{array} \right\}. \quad (9)$$

Similarly, for short-term sex, either for infected or treated individuals ($\phi = 0, t$), the analog to (2) is

$$\begin{aligned} \tilde{V}_s^\iota(\phi, x) = & \ln(y - z_s) + pI(s) + u[1 - I(s)] + \alpha_\phi \iota A \\ & + [1 - \alpha_\phi] \iota \left\{ \begin{array}{l} Q(\phi)[\eta V_t^\beta(t, x) + (1 - \eta)V_t^\iota(t, x)] \\ + [1 - Q(\phi)][\eta V_t^\beta(\phi, x) + (1 - \eta)V_t^\iota(\phi, x)] \end{array} \right\}, \quad (10) \end{aligned}$$

for $s = p, u$. For young healthy individuals ($\phi = 1$) the analog to (3) is

$$\begin{aligned} \tilde{V}_s^\iota(1, x) = & \ln(y - z_s) + pI(s) + u[1 - I(s)] \\ & + \sum_{\hat{\phi}} R_s(\hat{\phi}) [1 - \gamma_s(\hat{\phi})] \chi(c) \iota \left\{ \begin{array}{l} q[\eta V_t^\beta(t, x) + (1 - \eta)V_t^\iota(t, x)] \\ + (1 - q)[\eta V_t^\beta(0, x) + (1 - \eta)V_t^\iota(0, x)] \end{array} \right\} \\ & + \{1 - \sum_{\hat{\phi}} R_s(\hat{\phi}) [1 - \gamma_s(\hat{\phi})] \chi(c)\} \iota [\eta V_t^\beta(1, x) + (1 - \eta)V_t^\iota(1, x)]. \end{aligned} \quad (11)$$

For long-term sex, note that the transition probabilities $\Upsilon(\phi', \hat{\phi}' | \phi, \hat{\phi}, c, \hat{c})$ in Appendix D.2 are not affected by the discount factor, and therefore the young indi-

vidual's analog of (6) is

$$\begin{aligned}
\tilde{V}_l^\iota(\phi, \hat{\phi}, \hat{c}, x) &= \ln(y - z_l) + u + l + \alpha_\phi \iota A \\
&+ (1 - \alpha_\phi)(1 - \epsilon)(1 - \delta)(1 - \alpha_{\hat{\phi}})\iota \times \\
&\sum_{\phi', \hat{\phi}'} \Upsilon(\phi', \hat{\phi}' | \phi, \hat{\phi}, c, \hat{c}) \left[\begin{array}{c} \eta \tilde{V}_l^\beta(\phi', \hat{\phi}', x) \\ +(1 - \eta) \tilde{V}_l^\iota(\phi', \hat{\phi}', x) \end{array} \right] \\
&+ (1 - \alpha_\phi)[1 - (1 - \epsilon)(1 - \delta)(1 - \alpha_{\hat{\phi}})]\iota \times \\
&\sum_{\phi', \hat{\phi}'} \Upsilon(\phi', \hat{\phi}' | \phi, \hat{\phi}, c, \hat{c}) \left[\begin{array}{c} \eta V_l^\beta(\phi', x) \\ +(1 - \eta) V_l^\iota(\phi', x) \end{array} \right]. \tag{12}
\end{aligned}$$

D.2 Transition Probabilities in Long-term Relationships

The transition probabilities, $\Upsilon(\phi', \hat{\phi}' | \phi, \hat{\phi}, c, \hat{c})$, from the situation where a relationship is currently characterized by the quadruple $(\phi, \hat{\phi}, c, \hat{c})$ to the couple's health status next period or to $(\phi', \hat{\phi}')$ are now presented. Start with the situation where the person is currently healthy ($\phi = 1$) but his partner is infected ($\hat{\phi} \in \{0, t\}$). The following lists all possible cases for this situation:

$$\begin{aligned}
\Upsilon(1, t | 1, \hat{\phi}, c, \hat{c}) &= \{1 - [1 - \gamma_u(\hat{\phi})]\chi(c)\}Q(\hat{\phi}); \\
\Upsilon(1, 0 | 1, \hat{\phi}, c, \hat{c}) &= \{1 - [1 - \gamma_u(\hat{\phi})]\chi(c)\}[1 - Q(\hat{\phi})]; \\
\Upsilon(0, t | 1, \hat{\phi}, c, \hat{c}) &= [1 - \gamma_u(\hat{\phi})]\chi(c)(1 - q)Q(\hat{\phi}); \\
\Upsilon(0, 0 | 1, \hat{\phi}, c, \hat{c}) &= [1 - \gamma_u(\hat{\phi})]\chi(c)(1 - q)[1 - Q(\hat{\phi})]; \\
\Upsilon(t, t | 1, \hat{\phi}, c, \hat{c}) &= [1 - \gamma_u(\hat{\phi})]\chi(c)qQ(\hat{\phi}); \\
\Upsilon(t, 0 | 1, \hat{\phi}, c, \hat{c}) &= [1 - \gamma_u(\hat{\phi})]\chi(c)q[1 - Q(\hat{\phi})];
\end{aligned} \tag{13}$$

The chance that the individual remains healthy is given by $\{1 - [1 - \gamma_u(\hat{\phi})]\chi(c)\}$, while the odds that they will not are $[1 - \gamma_u(\hat{\phi})]\chi(c)$. In the latter case, the person will get treated with probability q and not with $1 - q$. The term $Q(\hat{\phi})$ reflects the odds of the partner being treated while the one $1 - Q(\hat{\phi})$ gives the odds that the companion is not.

The symmetric probabilities for when the partner is healthy ($\hat{\phi} = 1$) but the indi-

vidual is infected or treated ($\widehat{\phi} = 1$) are given below:

$$\begin{aligned}
\Upsilon(t, 1|\phi, 1, \widehat{c}, c) &= \{1 - [1 - \widehat{\gamma}_u(\phi)]\chi(\widehat{c})\}Q(\phi); \\
\Upsilon(0, 1|\phi, 1, \widehat{c}, c) &= \{1 - [1 - \widehat{\gamma}_u(\phi)]\chi(\widehat{c})\}[1 - Q(\phi)]; \\
\Upsilon(t, 0|\phi, 1, \widehat{c}, c) &= [1 - \widehat{\gamma}_u(\phi)]\chi(\widehat{c})(1 - q)Q(\phi); \\
\Upsilon(0, 0|\phi, 1, \widehat{c}, c) &= [1 - \widehat{\gamma}_u(\phi)]\chi(\widehat{c})(1 - q)[1 - Q(\phi)]; \\
\Upsilon(t, t|\phi, 1, \widehat{c}, c) &= [1 - \widehat{\gamma}_u(\phi)]\chi(\widehat{c})qQ(\phi); \\
\Upsilon(0, t|\phi, 1, \widehat{c}, c) &= [1 - \widehat{\gamma}_u(\phi)]\chi(\widehat{c})q[1 - Q(\phi)].
\end{aligned} \tag{14}$$

Next, both partners might be infected ($\phi \in \{0, t\}$ and $\widehat{\phi} \in \{0, t\}$), in which case a healthy future is no longer an option. The only question that remains is whether the future sees treatment or not, so that

$$\Upsilon(\phi', \widehat{\phi}'|\phi, \widehat{\phi}, c, \widehat{c}) = \begin{cases} [1 - Q(\phi)]Q(\widehat{\phi}), & \text{for } (\phi', \widehat{\phi}') = (0, t); \\ [1 - Q(\phi)][1 - Q(\widehat{\phi})], & \text{for } (\phi', \widehat{\phi}') = (0, 0); \\ Q(\phi)Q(\widehat{\phi}), & \text{for } (\phi', \widehat{\phi}') = (t, t); \\ Q(\phi)[1 - Q(\widehat{\phi})], & \text{for } (\phi', \widehat{\phi}') = (t, 0). \end{cases} \tag{15}$$

The last remaining case is where both partners are currently healthy. Here, $\Upsilon(1, 1|1, 1, c, \widehat{c}) = 1$, implying that $\Upsilon(\phi', \widehat{\phi}'|1, 1, c, \widehat{c}) = 0$, when $\phi' \in \{0, t\}$ and/or $\widehat{\phi}' \in \{0, t\}$.

D.3 Stationary Equilibrium

A stationary equilibrium for the developed framework will now be formulated. First, the equilibrium distributions for singles will be specified. Let $\mathcal{S}^d(\phi; x)$ represent the (non-normalized) stationary distribution of singles at the beginning of a period. It denotes the measure of singles of type- x that have status ϕ and discount factor d . Similarly, let $\mathcal{L}^d(\phi, \widehat{\phi}; x, \widehat{x})$ denote the measure of long-term relationships with status ϕ and $\widehat{\phi}$ and types x and \widehat{x} . Given some distributions \mathcal{S} and \mathcal{L} of singles and married people, the sexual behavior of individuals according to their decision rule Π [$\Pi_{g,r}^d = \Pi_{g,r}^d(\phi, x)$ for each status and type] gives rise to a new distribution of singles and married people, which can be described by a mapping \mathbf{T} that is characterized fully in Section D.4. In steady-state the distributions of singles and married people remain constant, and are determined by a

fixed point of this operator:

$$(\mathcal{S}^\beta, \mathcal{L}^\beta, \mathcal{S}^\iota, \mathcal{L}^\iota) = \mathbf{T}(\mathcal{S}^\beta, \mathcal{L}^\beta, \mathcal{S}^\iota, \mathcal{L}^\iota; \Pi). \quad (16)$$

Next, the expectations over the fraction of types in each market have to be consistent with the aggregation of individual choices in equilibrium. It is now useful to introduce the subscript g (for $g = f, m$) to a function or variable to denote the gender of the person in question. The number of market participants for sexual activity r ($= l, p, u$), who are of gender g and type- x with status ϕ and discount factor d , is given by

$$\mathcal{M}_{g,r}^d(\phi, x) \equiv \begin{cases} \Pi_{g,r}^d(\phi, x) \mathcal{S}_g^d(\phi; x), & \text{if } r = l, \\ [1 - \Pi_{g,l}^d(\phi, x)] \Pi_{g,r}^d(\phi, x) \mathcal{S}_g^d(\phi; x), & \text{if } r = p, u. \end{cases} \quad (17)$$

The number of market participants equals the number of singles times their probability of participating in a particular market. For the short-term market this also entails the probability of not previously finding a long-term partner within the current period. Then the fraction $R_{s,r}(\phi)$ of agents of status ϕ in market s of gender g is given by

$$R_{g,s}(\phi) = \frac{\sum_d \sum_x \mathcal{M}_{g,s}^d(\phi, x)}{\sum_d \sum_x \sum_{\phi'} \mathcal{M}_{g,s}^d(\phi', x)}, \text{ for all } g \text{ and } s \in \{p, u\}. \quad (18)$$

For the long-term market, the relevant fraction is given by:

$$R_{g,l}(\phi, c) = \frac{\sum_d \sum_x \mathcal{M}_{g,l}^d(\phi, x) \mathcal{I}(c(x) = c)}{\sum_d \sum_x \sum_{\phi'} \mathcal{M}_{g,l}^d(\phi', x)}, \text{ for all } g, \quad (19)$$

where $c(x)$ is a slight abuse of notation that denotes the circumcision status of the agent that is contained in his or her type x . The function $\mathcal{I}(\cdot)$ is an indicator function that takes the value of 1 if its argument is true. Note that $R_{f,s}(\phi)$ and $R_{f,l}(\phi, c)$ denote the distributions among women, which are relevant for *men* when determining their odds of getting infected. Similarly, $R_{m,s}(\phi)$ and $R_{m,l}(\phi, c)$ refer to the odds among men, but are relevant for the *women* when making their decisions.

Market clearing requires that the number of female participants equals the number of male participants in any market:

$$\sum_d \sum_x \sum_\phi \mathcal{M}_{f,r}^d(\phi, x) = \sum_d \sum_x \sum_\phi \mathcal{M}_{m,r}^d(\phi, x), \text{ for all } r. \quad (20)$$

Additionally, a transfer paid by one gender on a market is a transfer earned by the other so that

$$z_{f,r} + z_{m,r} = 0, \text{ for all } r. \quad (21)$$

This leads to the following formal definition of equilibrium.

Definition. A stationary equilibrium is described by a set of decision rules for search effort, $\Pi_{g,r}^d(\phi, x)$, a set of transfer payments, $z_{g,r}$, a set of stationary distributions, $\mathcal{S}_g^d(\phi; x)$ and $\mathcal{L}^d(\phi, \hat{\phi}; x, \hat{x})$, and status/type prevalence in each market, $R_{g,s}(\phi)$ and $R_{g,l}(\phi, c)$, for all $d = \{\iota, \beta\}$, $g \in \{f, m\}$, $r \in \{l, p, u\}$, $s \in \{p, u\}$, such that:

1. The decision rules for search intensities, $\Pi_{g,r}^d(\phi, x)$, satisfy the appropriately gender subscripted versions of the generic problems (4) and (8), taking as given transfer payments and HIV/AIDS prevalence rates;
2. The stationary distributions, $\mathcal{S}_g^d(\phi; x)$ and $\mathcal{L}^d(\phi, \hat{\phi}; x, \hat{x})$, solve the appropriately gender subscripted version of (16);
3. The distributions over health status for each market, $R_{g,s}(\phi)$ and $R_{g,l}(\phi, c)$, are given by (18) and (19) using (17);
4. The transfer payments, $z_{r,g}$, are such that the markets for all types of relationships clear according to (20). Additionally, the flow of transfers across the genders must balance as specified by (21).

D.4 Stationary Distributions

The transition operator \mathbf{T} defined in the Section D.3 is now fully characterized. Before starting, recall the function $\mathcal{I}(\cdot)$, an indicator function that takes the value of 1 if its argument is true and 0 otherwise. Focus on a particular gender so that

the gender subscript can be omitted. At the beginning of a period, recall that $\mathcal{S}^d(\phi; x)$ denotes the mass of singles with discount factor d , health status ϕ , and type x . $\mathcal{L}^d(\phi, \hat{\phi}; x, \hat{x})$ represents the mass of married individuals of HIV status ϕ of type x , discount factor d , and partners with HIV status $\hat{\phi}$ of type \hat{x} . Let \mathcal{A} be the mass of individuals with the final symptoms of AIDS. The sexual behavior of individuals is governed by their decision rules, $\pi_{g,r}^d = \Pi_{g,r}^d(\phi, x)$.

Assume temporarily that only people who are of health status $\phi = t$ will be treated next period. Suppose also that the newborns continue on with a low discount factor in the next period. Moreover, assume the individual's discount factor does not change. Given the beginning of period distributions \mathcal{S}^d , \mathcal{L}^d , and \mathcal{A} one can compute the distributions at the beginning of next period under these assumptions. Call these \mathcal{S}'^d , \mathcal{L}'^d , and \mathcal{A}' . These can then be adjusted for changing treatment status and discount factors.

For notational convenience, define $\tilde{Q}(\hat{\phi})$ to be such that $\tilde{Q}(0) = q$ and $\tilde{Q}(t) = q$ (this is different from Q defined in the main text). Moreover, define the following variables to represent the infectiousness of each short-term market:

$$\hat{\theta}_s = \sum_{\hat{\phi}} R_s(\hat{\phi})[1 - \gamma_s(\hat{\phi})], \text{ for all } s \in \{p, u\}. \quad (22)$$

First consider *next period's distribution of single individuals*. Consider first the distribution of healthy singles next period:

$$\begin{aligned} \mathcal{S}'^d(1, x) &= (1 - \delta) \times \{ \\ &\quad \mathcal{S}^d(1, x)[1 - \Pi_l^d(1, x)]\{1 - \Pi_p^d(1, x) - \Pi_u^d(1, x) + \sum_s \Pi_s^d(1, x)[1 - \hat{\theta}_s \chi(c)]\} \\ &\quad + \sum_{\hat{\phi}, \hat{x}} [\mathcal{L}^d(1, \hat{\phi}; x, \hat{x}) + R_l(\hat{\phi}, \hat{c})\Pi_l^d(1, x)\mathcal{S}^d(1, x)]\{1 - [1 - \gamma_u(\hat{\phi})]\chi(c)\} \\ &\quad \times [1 - (1 - \delta)(1 - \alpha_{\hat{\phi}})(1 - \varepsilon)] \} \\ &\quad + \mu(x)\mathcal{I}(d = \iota). \end{aligned} \quad (23)$$

Singles survive with probability $(1 - \delta)$, captured by the first line. The second line accounts for healthy singles this period that continue as healthy singles. There are $\mathcal{S}^d(1, x)$ such singles this period. They remain healthy singles if they do not

successfully enter the long-term market represented by the first bracket and if they either do not enter the short-term market or enter but do not get infected, as represented by the braces. The third line accounts for those who return as healthy singles from marriage. The bracket gives the stock of individuals married to a partner of status $\hat{\phi}$ at the start of the period plus those singles who newly marry such a partner this period. They remain healthy with probability $\{1 - [1 - \gamma_u(\hat{\phi})]\chi(c)\}$, but the marriage breaks up with the probability in the bracket on the fourth line. The final line accounts for the newborns.

Consider next the distribution of infected individuals without treatment next period:

$$\begin{aligned}
\mathcal{S}^{d}(0, x) &= (1 - \delta) \times \{ \\
&\quad \mathcal{S}^d(1, x)[1 - \Pi_l^d(1, x)] \sum_s \Pi_s^d(1, x) \hat{\theta}_s \chi(c) \\
&\quad + \sum_{\hat{\phi}, \hat{x}} [\mathcal{L}(1, \hat{\phi}; x, \hat{x}) + R_l(\hat{\phi}, \hat{c}) \Pi_l^d(1, x) \mathcal{S}^d(1, x)] [1 - \gamma_u(\hat{\phi})] \chi(c) \\
&\quad \times [1 - (1 - \delta)(1 - \alpha_{\hat{\phi}})(1 - \varepsilon)] \\
&\quad + \mathcal{S}^d(0, x)(1 - \alpha_0)[1 - \Pi_l^d(0, x)] \\
&\quad + (1 - \alpha_0) \sum_{\hat{\phi}, \hat{x}} [\mathcal{L}^d(0, \hat{\phi}; x, \hat{x}) + R_l(\hat{\phi}, \hat{c}) \Pi_l^d(0, x) \mathcal{S}^d(0, x)] \\
&\quad \times [1 - (1 - \delta)(1 - \alpha_{\hat{\phi}})(1 - \varepsilon)] \} \tag{24}
\end{aligned}$$

The first four lines capture the same elements as in the previous equation, but now healthy individuals only transit to state $\phi = 0$ if they get infected. Additionally, now individuals that are already infected may survive to the next period. Line five captures single infected individuals, who do not develop final stage symptoms with probability $1 - \alpha_0$ and do not enter the long-term market with probability $1 - \Pi_l^d(0, x)$, and therefore survive as infected singles. Lines six and seven capture individuals that either started in marriage or got married, similar to lines three and four, but now these individuals are infected. Again they return as infected singles, if they do not develop final stage symptoms, and if the marriage does not survive.

Finally, the distribution of treated individuals next period is given by

$$\begin{aligned}
\mathcal{S}'^d(t, x) = & (1 - \delta) \times \{ \\
& (1 - \alpha_t) \mathcal{S}^d(t, x) [1 - \Pi_l^d(t, x)] \\
& + (1 - \alpha_t) \sum_{\hat{\phi}, \hat{x}} [\mathcal{L}^d(t, \hat{\phi}; x, \hat{x}) + R_l(\hat{\phi}, \hat{c}) \Pi_l^d(t, x) \mathcal{S}^d(t, x)] \\
& \times [1 - (1 - \delta)(1 - \alpha_{\hat{\phi}})(1 - \varepsilon)] \} \tag{25}
\end{aligned}$$

The four lines here correspond to lines one, five, six and seven in the previous expression. The reason the intermediate lines are dropped is the temporary assumption that only individuals who were already in treatment at the beginning of the period are eligible for treatment next period. This will be adjusted later.

The mass of individuals with final stage symptoms next period is

$$\mathcal{A} = \sum_{d, x} (1 - \delta) \{ (1 - \delta_2) \mathcal{A} + [\mathcal{S}^d(\phi, x) + \sum_{\hat{\phi}, \hat{x}} \mathcal{L}^d(\phi, \hat{\phi}, x, \hat{x})] \alpha_{\phi} \}. \tag{26}$$

It comprises those that started the period in the final stage and neither died of natural causes nor of AIDS related reasons. It also includes all other individuals who develop final stage symptoms, which occurs with probability α_{ϕ} .

Now consider the *distribution of long-term marriages next period*. Consider a marriage with two healthy individuals. The marriage survives if neither spouse dies of natural causes and the marriage does not break up exogenously. The stock of marriages next period includes marriages in current period made up both from previous ones plus new ones. The mass of marriages next period is given by

$$\begin{aligned}
\mathcal{L}'^d(1, 1; x, \hat{x}) = & (1 - \delta)^2 (1 - \varepsilon) \\
& \times [\mathcal{L}^d(1, 1; x, \hat{x}) + (1 - R_l(0, \hat{c}) - R_l(t, \hat{c})) \Pi_l^d(1, x) \mathcal{S}^d(1, x)]. \tag{27}
\end{aligned}$$

Consider marriages where the partner is infected or treated. The terms are similar to before, only that now the marriage breaks up for one additional reason: if the partner develops symptoms (probability $\alpha_{\hat{\phi}}$). The person also stays healthy with probability $\{1 - [1 - \gamma_u(\hat{\phi})] \chi(c)\}$. So,

$$\begin{aligned}\mathcal{L}'^d(1, \widehat{\phi}; x, \widehat{x}) &= (1 - \delta)^2(1 - \varepsilon)(1 - \alpha_{\widehat{\phi}})\{1 - [1 - \gamma_u(\widehat{\phi})]\chi(c)\} \\ &\times [\mathcal{L}^d(1, \widehat{\phi}; x, \widehat{x}) + R_l(\widehat{\phi}, \widehat{c})\Pi_l^d(1, x)\mathcal{S}^d(1, x)].\end{aligned}\quad (28)$$

Similar expressions obtain for partnerships where the individual under consideration is infected or treated but the partner from the other gender is healthy:

$$\begin{aligned}\mathcal{L}'^d(\phi, 1; x, \widehat{x}) &= (1 - \delta)^2(1 - \varepsilon)(1 - \alpha_{\phi})\{1 - [1 - \widehat{\gamma}_u(\phi)]\chi(\widehat{c})\} \\ &\times [\mathcal{L}^d(\phi, 1; x, \widehat{x}) + R_l(1, \widehat{c})\Pi_l^d(\phi, x)\mathcal{S}^d(\phi, x)].\end{aligned}\quad (29)$$

In the case where both spouses are infected or treated, it is no longer required to take into account the transmission of the disease, but there is a chance that either one will develop symptoms:

$$\begin{aligned}\mathcal{L}'^d(\phi, \widehat{\phi}; x, \widehat{x}) &= (1 - \delta)^2(1 - \varepsilon)(1 - \alpha_{\phi})(1 - \alpha_{\widehat{\phi}}) \\ &\times [\mathcal{L}^d(\phi, \widehat{\phi}; x, \widehat{x}) + R_l(\widehat{\phi}, \widehat{c})\Pi_l^d(\phi, x)\mathcal{S}^d(\phi, x)].\end{aligned}\quad (30)$$

Finally, introduce the *adjustments for a changing discount factor and changing treatment status*, which are mechanical parts of the model that do not involve many choices. Incorporate discount factor changes first. To do this, let D''^d represent generic auxiliary distributions that result from incorporating the transitions from the previous D'^d due to changing discount factors. High-discount factor individuals stay with a high discount factor, but also a fraction η of low-discount factor individuals change to a high discount factor:

$$D''^\beta(\phi, \dots) = D'^\beta(\phi, \dots) + \eta D'^\mu(\phi, \dots) \quad (31)$$

$$D''^\mu(\phi, \dots) = (1 - \eta)D'^\mu(\phi, \dots). \quad (32)$$

To give examples, $\mathcal{S}''^\beta(\phi, x) = \mathcal{S}'^\beta(\phi, x) + \eta\mathcal{S}'^\mu(\phi, x)$ counts the number of type- x singles of health status ϕ that end up with the high discount, β . Another example would be $\mathcal{L}''^\mu(\phi, \widehat{\phi}; x, \widehat{x}) = (1 - \eta)\mathcal{L}'^\mu(\phi, \widehat{\phi}; x, \widehat{x})$.

The true distributions next period, which in steady state are the original distributions from the beginning and therefore with slight abuse of notation did not

receive a new name, are recovered by taking into account changes in treatment status, since infected individuals with status 0 change to a treated status t with probability q :

$$\begin{aligned}
\mathcal{S}^d(1, x) &= \mathcal{S}^{nd}(1, x), \\
\mathcal{S}^d(0, x) &= \mathcal{S}^{nd}(0, x)(1 - q), \\
\mathcal{S}^d(t, x) &= \mathcal{S}^{nd}(0, x)q + \mathcal{S}^{nd}(t, x), \\
\mathcal{L}^d(1, 1, x, \hat{x}) &= \mathcal{L}^{nd}(1, 1, x, \hat{x}), \\
\mathcal{L}^d(1, 0, x, \hat{x}) &= \mathcal{L}^{nd}(1, 0, x, \hat{x})(1 - q), \\
\mathcal{L}^d(0, 1, x, \hat{x}) &= \mathcal{L}^{nd}(0, 1, x, \hat{x})(1 - q), \\
\mathcal{L}^d(0, 0, x, \hat{x}) &= \mathcal{L}^{nd}(0, 0, x, \hat{x})(1 - q)^2, \\
\mathcal{L}^d(1, t, x, \hat{x}) &= \mathcal{L}^{nd}(1, 0, x, \hat{x})q + \mathcal{L}^{nd}(1, t, x, \hat{x}), \\
\mathcal{L}^d(t, 1, x, \hat{x}) &= \mathcal{L}^{nd}(0, 1, x, \hat{x})q + \mathcal{L}^{nd}(t, 1, x, \hat{x}), \\
\mathcal{L}^d(0, t, x, \hat{x}) &= \mathcal{L}^{nd}(0, t, x, \hat{x})(1 - q) + \mathcal{L}^{nd}(0, 0, x, \hat{x})(1 - q)q, \\
\mathcal{L}^d(t, 0, x, \hat{x}) &= \mathcal{L}^{nd}(t, 0, x, \hat{x})(1 - q) + \mathcal{L}^{nd}(0, 0, x, \hat{x})q(1 - q), \\
\mathcal{L}^d(t, t, x, \hat{x}) &= \mathcal{L}^{nd}(t, t, x, \hat{x}) + \mathcal{L}^{nd}(0, t, x, \hat{x})q + \mathcal{L}^{nd}(t, 0, x, \hat{x})q + \mathcal{L}^{nd}(0, 0, x, \hat{x})q^2.
\end{aligned}$$

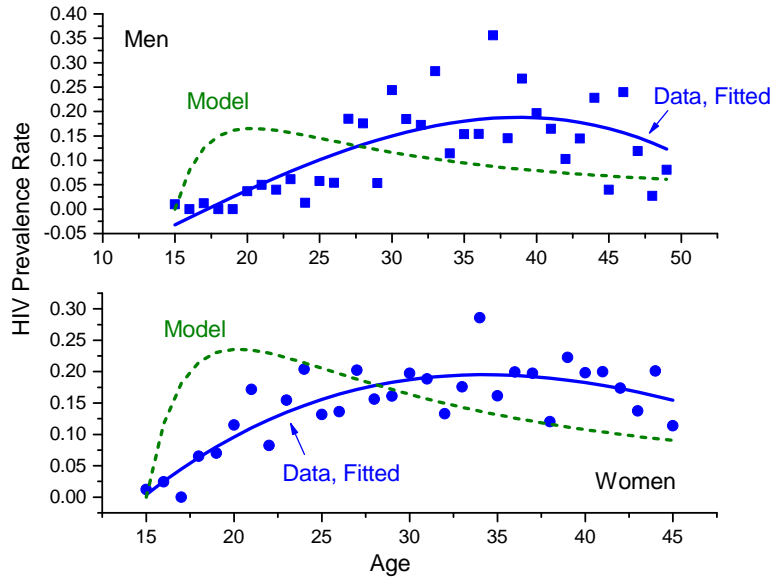
The right-hand sides of these equations together with (23) to (32) fully describe the operator \mathbf{T} in (16).

E Appendix—Life-cycle Implications

Figure 9 plots HIV/AIDS prevalence by age.⁵⁶ Both the model and data agree on a hump-shaped infection pattern, despite the fact that agents in the model become sexually active earlier than is observed in the data, which shifts the model's life-cycle predictions on HIV/AIDS infections to the left. The hump-shaped pattern is explained by two opposing forces. First, the rise in HIV/AIDS infection is due to the fact that older people have had more time to be sexually active, and so a larger percentage of the older population is infected with HIV/AIDS. Second,

⁵⁶The data is fitted with a third-order polynomial. The somewhat choppy raw data is due to the small sample sizes.

Figure 9: HIV Prevalence Rate—Men vs. Women, Model vs. Data



people who are infected early in life will die before they make it to old age. Put differently, people who have made it to old age must be those who have engaged in less risky sexual behavior and so are less likely to be infected with HIV/AIDS. This second effect explains the eventual drop in HIV/AIDS prevalence seen at older ages. Figure 9 also illustrates the differential patterns of infection between the sexes. The figure shows that women get infected earlier than men, both in the model and the data.

The model also does a very nice job in matching the decline in risky activity over the life cycle. Older people are less likely to be single, see Figure 10. As people age, they are thus less likely to engage in casual sex; this is also reported in Figure 10. The fact that the discount factor stochastically rises with age helps to generate this pattern.

An additional prediction of the model relates to death causes, since individuals may die due to HIV/AIDS or due to other natural causes. Figure 11 compares the model prediction over the life cycle with its data counterpart. Both the model and the data exhibit a hump-shaped pattern of HIV/AIDS-caused deaths; this is consistent with the hump-shaped pattern of infection rates.

Figure 10: Sexual Behavior by Age—Model vs. Data

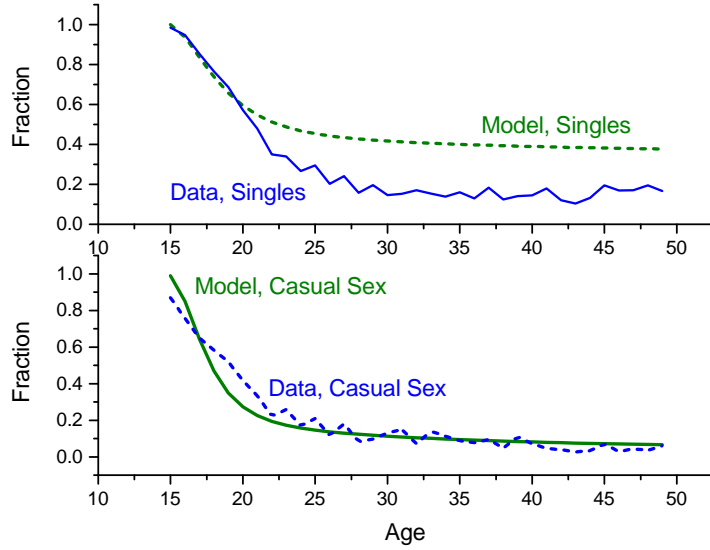
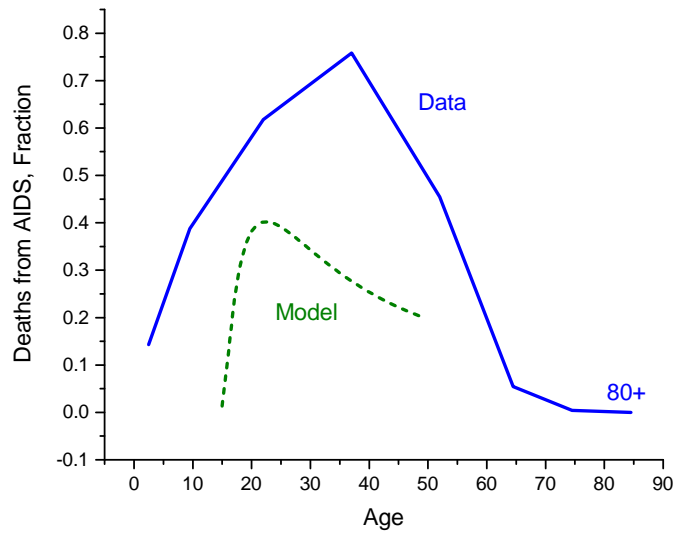


Figure 11: Deaths by HIV /AIDS by Age, Fraction—Model vs. Data



F Appendix—More Details on Model Extensions

F.1 The Diffusion of Better Information—Details

To capture uninformed individuals, consider a type- x person with $i = 0$. Let $\mathring{s} \in \{a, p, u\}$ and define $\mathring{I}(\mathring{s}) = 1$, for $\mathring{s} = p$, and $\mathring{I}(\mathring{s}) = 0$, otherwise. Likewise, define $J_{\mathring{s}}(\mathring{s}) = 0$, if $\mathring{s} = a$, and $J_{\mathring{s}}(\mathring{s}) = 1$, otherwise. If healthy, the value from short-term sex, for $\mathring{s} = \{a, u, p\}$, is now given by

$$\begin{aligned} \tilde{V}_{\mathring{s}}^{\beta}(1, x) &= \ln(y - z_{\mathring{s}}) + \{p\mathring{I}(\mathring{s}) + u[1 - \mathring{I}(\mathring{s})]\}J_{\mathring{s}}(\mathring{s}) \\ &\quad + \sum_{\hat{\phi}} R_u(\hat{\phi})[1 - \gamma_u(\hat{\phi})]\beta[qV_i^{\beta}(t, x) + (1 - q)V_i^{\beta}(0, x)] \\ &\quad + \{1 - \sum_{\hat{\phi}} R_u(\hat{\phi})[1 - \gamma_u(\hat{\phi})]\}\beta V_i^{\beta}(1, x), \end{aligned} \quad (33)$$

where $z_a = 0$. Compared to the value function for informed individuals (3), uninformed people perceive all sex as being as risky as unprotected sex without circumcision. In the case of abstinence, the uninformed now worry about infection even when they don't have sex — cf. (1).

In the long-term market an uninformed individual thinks that transmissions are governed as if people are not circumcised. That means that for a uninformed type- x individual (so that $i = 0$)

$$\begin{aligned} \tilde{V}_i^{\beta}(\phi, \hat{\phi}, \hat{c}, x) &= \ln(y - z_l) + u + l + \alpha_{\phi}\beta A \\ &\quad + (1 - \alpha_{\phi})(1 - \epsilon)(1 - \delta)(1 - \alpha_{\hat{\phi}})\beta \sum_{\phi', \hat{\phi}'} \Upsilon(\phi', \hat{\phi}' | \phi, \hat{\phi}, 0, 0) \tilde{V}_i^{\beta}(\phi', \hat{\phi}', \hat{c}, x) \\ &\quad + (1 - \alpha_{\phi})[1 - (1 - \epsilon)(1 - \delta)(1 - \alpha_{\hat{\phi}})]\beta \sum_{\phi', \hat{\phi}'} \Upsilon(\phi', \hat{\phi}' | \phi, \hat{\phi}, 0, 0) V_i^{\beta}(\phi', x). \end{aligned}$$

(Note that c and \hat{c} have been set to 0 in the transition probability Υ .) Similar

adjustments need to be made for young uninformed type- x individuals. Now,

$$\begin{aligned}\tilde{V}_s^l(1, x) &= \ln(y - z_s) + \{p\dot{I}(\dot{s}) + u[1 - \dot{I}(\dot{s})]\}J_s(\dot{s}) \\ &+ \sum_{\hat{\phi}} R_u(\hat{\phi})[1 - \gamma_u(\hat{\phi})] \left\{ \begin{array}{l} q[\eta V_l^\beta(t, x) + (1 - \eta)V_l^l(t, x)] \\ +(1 - q)[\eta V_l^\beta(0, x) + (1 - \eta)V_l^l(0, x)] \end{array} \right\} \\ &+ \{1 - \sum_{\hat{\phi}} R_u(\hat{\phi})[1 - \gamma_u(\hat{\phi})]\} \eta V_l^\beta(1, x) + (1 - \eta)V_l^l(1, x),\end{aligned}$$

for $\dot{s} = \{a, u, p\}$. Last, the value function for a young uninformed individual in a long-term relationship is

$$\begin{aligned}\tilde{V}_l^l(\phi, \hat{\phi}, \hat{c}, x) &= \ln(y - z_l) + u + l + \alpha_\phi l A \\ &+ (1 - \alpha_\phi)(1 - \epsilon)(1 - \delta)(1 - \alpha_{\hat{\phi}}) \eta \times \\ &\sum_{\phi', \hat{\phi}'} \Upsilon(\phi', \hat{\phi}' | \phi, \hat{\phi}, 0, 0) \left[\begin{array}{l} \eta \tilde{V}_l^\beta(\phi', \hat{\phi}', x) \\ +(1 - \eta) \tilde{V}_l^l(\phi', \hat{\phi}', x) \end{array} \right] \\ &+ (1 - \alpha_\phi)[1 - (1 - \epsilon)(1 - \delta)(1 - \alpha_{\hat{\phi}})] \eta \times \\ &\sum_{\phi', \hat{\phi}'} \Upsilon(\phi', \hat{\phi}' | \phi, \hat{\phi}, 0, 0) \left[\begin{array}{l} \eta V_l^\beta(\phi', x) \\ +(1 - \eta) V_l^l(\phi', x) \end{array} \right].\end{aligned}$$

Table 13 reports the results of selected policy experiments using the environment with uninformed individual as a starting point. Such results are on the panel “1996 benchmark”. For ease of comparison, the table also reports the results using the baseline calibration (“2004 benchmark” panel). As discussed in Section 6, in general, policies implemented when there are uninformed agents (“1996 benchmark”) tend to have stronger effects. The reason is that uninformed agents do not change their behavior in response to the policy change.

F.2 Fertility—Details

This little appendix describes how the model is parameterized when there is heterogeneous utility for fertility within marriage, as described in Section 5.3.1. Note that, conditional on an individual’s circumcision status, there is only one source of permanent heterogeneity in the benchmark model: the time discount factors, β

Table 13: IMPROVED INFORMATION—MODEL VS. DATA

	Bench.	Condoms $p = 5.5$	Circumcision 100%	ART $q = 0.01$	ART $q = 0.075$	STDs $\gamma_u^m = 0.9$
Panel (a): 1996 benchmark						
HIV prevalence, %	14.4	20.7	4.5	12.7	7.0	11.1
Sex that is casual, %	19.8	34.6	26.2	21.6	35.0	20.2
Condom use (singles), %	27.9	48.8	25.8	26.4	26.0	26.4
Fraction of singles, %	47.4	58.9	50.3	47.5	58.1	46.2
Non-abstinent singles, %	66.8	72.2	82.1	73.1	86.1	71.7
Panel (b): 2004 benchmark						
HIV prevalence, %	10.3	15.8	5.6	11.3	8.1	9.5
Sex that is casual, %	15.7	34.0	28.9	20.5	41.5	18.7
Condom use (singles), %	32.8	59.2	22.4	27.2	24.1	27.6
Fraction of singles, %	48.0	62.4	56.7	51.1	65.7	49.6
Non-abstinent singles, %	53.6	66.4	75.6	62.5	83.6	60.3

and ι . Since there is a one-to-one correspondence between β and ι , think about β as summarizing this heterogeneity. To keep the computational complexity fixed, leave the number of types constant. That is, for each type- x person with a different discount factor, β , assign a different level of utility from fertility $f(x)$; again, $f(x)$ is really just changing with β . Now, the instantaneous utility from a long-term relationship is given by $u + l + f(x)$.

Assume a quadratic form for the utility from fertility such that $f(x) = \theta_0 + \theta_1\beta + \theta_2\beta^2$, regardless of the value for c and ι . Note that θ_0 and l cannot be separately identified given the linear form of utility. Thus, three values must be calibrated: $\theta_0 + l$, θ_1 , and θ_2 . These are picked such that the new model fits the baseline data targets as closely as possible. The resulting parameters are: $\theta_0 = 28207$, $\theta_1 = -58380$ and $\theta_2 = 30201$. The fertility benefit $f(x)$ ranges from -1.55 to 21.23. Table 14 reports results from this new parameterization.

G Appendix—Computational Details

A capsule summary of the numerical algorithm used to solve the benchmark model is provided here. There are two key steps. The first step involves solving the model for a given set of parameter values. In the second step, the algorithm

TABLE 14: CALIBRATION—FERTILITY

<i>Observation</i>	<i>Data</i>	<i>Model benchmark</i>	<i>Model w/ fertility</i>
HIV prevalence, %	11.8	10.3	10.0
—Males	10	8.6	8.3
—Females	13	12.1	11.8
Sex that is casual, % (of all)	18	16	18
Condom use for casual sex, %	39	33	29
Singles that had casual sex in past year, %	47	54	65
Singles, %	33	48	47
Married by age 22, %			
—Males	58	57	54
—Females	90	63	58
Married by age 50, %			
—Males	100	98	78
—Females	100	98	88
Deaths related to HIV, %	29	25	32

picks the parameter values to match the model’s output with the data targets as closely as possible. The first step proceeds as follows:

1. The static problems (4) and (8) that yield the meeting probabilities are solved. The solution to these problems imply that each π can be implicitly written as a nonlinear function of the difference between either two \tilde{V} ’s or a \tilde{V} and V . Given a grid of values for the \tilde{V} ’s and V ’s, the π ’s can be computed using the bisection method. The values for the π ’s when the \tilde{V} ’s and V ’s lie off the grid can be obtained using an interpolation scheme.
2. One outer loop solves for the market-clearing prices using the NEWUOA algorithm. This algorithm picks the prices to minimize excess demand in the three relationship markets.
3. In an inner loop the value functions and stationary distributions are determined computationally, given prices, using standard iterative procedures. First the “matched” value functions (the \tilde{V} ’s) are computed for each type of individual. Then, the ex ante value functions (the V ’s) are calculated using a linear interpolation scheme that employs the results from 1. The stationary distributions are computed using the formulas in Appendices D.3 and

D.4.

In the second step the parameters values are calibrated using a Pattern Search algorithm. The calibration algorithm and the solutions to the static problems in 1 are implemented in MATLAB, while the more computationally demanding loops in 2 and 3 are coded in FORTRAN.