Health, Risky Behavior and the Value of Medical Innovation for Infectious Disease*

Tat Y. Chan† Barton H. Hamilton‡ Nicholas W. Papageorge§

September 15, 2015

Abstract: We propose a dynamic framework to study the value of medical innovation in the context of infectious disease. We apply our framework to evaluate an HIV treatment breakthrough known as HAART. The model captures how, in lowering both the expected cost and likelihood of HIV infection, HAART reduced the implicit price of risky sex. Forward-looking agents responded by optimally shifting their behavior. The model also imposes equilibrium constraints, explicitly capturing how optimal shifts in behavior affect equilibrium choices by changing both infection probabilities and the ease of finding partners willing to engage in risky sex. Using the estimated model, we conduct counterfactual simulations to compute the value of HAART from the perspective of uninfected agents. This includes the option value of the innovation along with value accruing from changes in sex behavior in response to HAART introduction. We also calculate the added-value of a fully functional vaccine from the perspective of both infected and uninfected agents, where infected agents benefit from a vaccine due to resulting shifts in market equilibrium.

Keywords: Innovation, Option Value, Equilibrium, HIV/AIDS, Risky Sex, Structural Models

JEL Classification: I1 J2

*We gratefully acknowledge helpful comments from: Scott Cunningham, Gautam Gowrisankaran, Mark Gersovitz, Maartin Lindeboom, Glenn MacDonald, Robert Miller, Robert Moffitt, Juan Pantano, Áureo de Paula, Robert Pollak, Carol Propper, John Rust, Maher Said, Seth Sanders, Paul Scott and Petra Todd and seminar participants at Washington University in Saint Louis, Johns Hopkins University, ITAM, Toulouse School of Economics, University of Bristol, City University of London, Imperial College London, Baylor University, the 2012 Annual Health Econometrics Workshop and the 2013 Economic Demography Workshop along with excellent research assistance from Andrew Gray, Shaiza Qayyum and Matthew White. We also thank the editor and five anonymous referees for their helpful suggestions. The usual caveats apply.

†Olin Business School, Washington University in St. Louis. chan@wustl.edu.

‡Olin Business School, Washington University in St. Louis. hamiltonb@wustl.edu [corresponding author].

§Department of Economics, Johns Hopkins University. papageorge@jhu.edu.

1 Introduction

Biomedical research expenditures, which in the U.S. exceed $100 billion annually, generate innovations upon existing medical treatments. Economic research evaluating these innovations tends to focus on the direct benefits to patients, including increases to longevity and improvements to the quality of life. In evaluating medical innovation, two critical factors are often overlooked. First, new medical technologies not only benefit individuals who are already ill, but also create an option value for forward-looking, healthy individuals who anticipate some risk of falling ill in the future. Second, by improving health trajectories conditional on illness, new medical technologies shift incentives and can therefore lead to changes in risky behavior. In the case of contagious disease, aggregate changes in risky behavior can lead to additional shifts in incentives, further affecting choices and the value of medical innovation. The reason is that infection rates for contagious disease are not only determined by individual risky behavior decisions, but are also influenced by how individuals interact with one another. Therefore, in the case of contagious disease, measuring the returns to biomedical research requires a framework that captures dynamic incentives along with how agents anticipate and respond to aggregate shifts in behavior.

In this paper, we propose a dynamic framework to measure the option value of medical innovation in the context of contagious disease. In the framework, forward-looking agents form expectations over the likelihood of infection along with the costs associated with being infected. The framework we develop hinges on the following idea: in many medical contexts, healthy individuals can influence their risk of illness by engaging in or refraining from risky behaviors.[1] In these contexts, medical breakthroughs reduce the implicit price of risky behavior by lowering the expected cost of illness. Forward-looking agents respond by optimally shifting their risky behavior. The magnitude of these shifts is informative, allowing us to identify preferences over risky behaviors, from which we can calculate the value of innovation.[2] To capture contagion, our framework also accounts for how agents form expectations over the behavior of other agents they may encounter. To that end, our dynamic decision-making framework embeds an equilibrium model that accounts for the optimal behavior changes of other agents in response to a medical innovation.

We apply our framework to analyze the sexual behavior of HIV-negative and HIV-positive

[1] Whereas some diseases are the consequence of factors (e.g., genetics) outside an individual's control, many others result from modifiable health behaviors on the part of the agent, such as smoking or physical inactivity. In fact, up to 50% of all deaths in the U.S. result from such risky health behaviors (McGinnis and Foege, 1993). Cawley and Ruhm (2012) provide an excellent summary of this literature.

[2] Similar to the idea in Thaler and Rosen (1976), who study occupational hazard, we value medical innovation by computing a compensation variation in terms of years of life.

men (hereafter: HIV− and HIV+, referring to those not infected and those infected with HIV, respectively) both before and after a medical innovation known as HAART that was introduced in 1996. HIV reduces the ability of the immune to system to fight off routine infections (a condition known as AIDS) and reached epidemic proportions in several countries, including the U.S., starting in 1984.[3] HIV and the AIDS epidemic provide a natural setting for studying the value of medical innovation from the perspective of healthy agents. First, the connection between risky sexual activity and HIV infection was both clearly established and actively disseminated shortly after the start of the epidemic. Individuals were therefore well aware of the potential consequences of their sexual decisions.[4] This is similar to other risky health behaviors such as smoking, where there is a clearly understood relationship to diseases such as lung cancer and emphysema.[5] The key dissimilarity is that the risk of sexual activity for an HIV− agent depends on the epidemic rate of HIV since infection through sex is only possible if the sexual partner is infected. Second, HIV infection is costly enough to affect agent decisions. Absent treatment, a newly infected HIV+ subject lives an average of 11 years and the advent of the AIDS epidemic brought drastic shifts in risky sexual behavior. Finally, HAART was an unexpected medical breakthrough that effectively transformed HIV infection from a virtual death sentence to a manageable chronic condition.[6] Within two years, the introduction of HAART reduced mortality rates by over 80% among HIV+ men (Bhaskaran et al., 2008). Furthermore, HAART dramatically reduced the rate of infection for agents whose partners are infected. Consequently, we are able to observe the response of healthy, forward-looking agents to a large, unanticipated, exogenous shift in the implicit costs of risky behavior. The introduction of HAART thus provides a quasi-experimental setting that allows us to identify preference parameters.

[3]AIDS stands for acquired immunodeficiency syndrome.

[4]Individuals in our sample are regularly tested for HIV and have periodic contact with medical specialists by virtue of being in the study. Therefore, the assumption of full information on HIV status, the consequences of risky behavior and availability of treatment is reasonable. Other work has considered contexts where this assumption is more problematic. Thornton (2008) studies the behavioral impact of learning HIV status and Paula, Shapira, and Todd (2014) study the effects of learning about the consequences of risky behavior, also in the context of HIV infection. Both studies consider a sample of individuals living in Malawi for whom information acquisition is likely to be more costly relative to individuals in our sample.

[5]Aggregate shifts in smoking behavior might have some feedback effect. For example, if enough people smoke more after the introduction of an innovation targeting lung cancer, one could imagine aggregate changes such as less social stigma associated with smoking, which could presumably enhance the value of the new medical technology among healthy individuals who enjoy cigarettes. Indeed, there is some evidence supporting the idea that stigma affects risky behaviors like smoking (Stuber, Galea, and Link, 2008).

[6]HAART stands for highly active anti-retroviral treatment. There is no vaccine or cure for HIV or AIDS, but HAART is the current standard treatment. In general, 1996 is marked as the year when two crucial clinical guidelines that comprise HAART came to be commonly acknowledged. First, protease inhibitors (made widely available towards the end of 1995) would be an effective HIV treatment. Second, several anti-retroviral drugs taken simultaneously could indefinitely delay the onset of AIDS.

We construct a dynamic choice model with several components. First, as HIV is a condition where individual behavior affects health outcomes, measuring the option value of innovations targeting HIV requires us to specify how individuals form expectations over the likelihood of illness conditional on their own behavior and continuation payoffs conditional on becoming ill. Further, HIV is an infectious disease where information is private. Therefore, the probability of infection is also influenced by the likelihood of being matched with an infectious partner.[7] This match probability is a function of market level factors, including the proportion of HIV+ men on the sex market at any given time, HIV+ men's infectiousness, which is partly determined by average usage of HAART, along with HIV+ men's optimal sex behavior. We also allow for the dynamic that if more HIV− agents choose risky sex once HAART is invented (a reaction known in the medical literature as *disinhibition*), HIV+ men will find it easier to find a willing partner for risky sex.[8] Taken together, an individual's optimal decisions are a function of the optimal behavior of other agents and equilibrium behavior occurs where HIV+ and HIV− men best respond to one another. In this sense, both conceptually and methodologically, our model is broadly similar to the equilibrium model of sex and matching in Arcidiacono, Beauchamp, and McElroy (2010), where teenage dating behavior is a function of preferences over partner characteristics along with the likelihood of matching with a given type, itself an output of the matching model capturing other agents' preferences.

The impact of dynamic considerations on current-period health behaviors has been studied in the context of drinking and smoking among the elderly (Arcidiacono, Sieg, and Sloan, 2007) and unprotected sex among teenaged females who anticipate how different safe-sex measures can affect the possibility of becoming pregnant (Arcidiacono, Khwaja, and Ouyang, 2012).[9] A key departure of the present study from earlier work is that we observe health-related decision-making both before and after a change in medical technology that drastically and abruptly shifted the consequences of risky behavior.[10] This exogenous shift is crucial in identifying preference parameters and also relates this project to previous work considering

[7]Estimating risky behavior in response to rates of infection is more similar to Ahituv, Hotz, and Philipson (1996), who study the demand for condoms as a function of local prevalence of AIDS.

[8]See Bechange et al. (2010), who study sexual disinhibition among HIV− individuals who are coupled with HIV+ individuals undergoing HIV treatment.

[9]Framing teenage unprotected sex in terms of rational decision-making is similar to our approach of investigating observed risky sex patterns as optimal choices. Other work applying economics to understand problematic behavior includes: Austen-Smith and Fryer (2005) and Fryer (2003), who construct theoretical models to rationalize seemingly poor educational decisions of black teenagers as best responses in a game played against their community.

[10]In a different context, Adda (2007) shows that behavior change can partially offset dangers associated with mad cow disease once consumers are informed of the risks.

the impact of price changes on health behaviors.[11]

This study also contributes to research on behavioral responses to the AIDS epidemic and HAART. Seminal work by Philipson and Posner (1993) and Kremer (1996) introduced behavioral choice into epidemiological models of AIDS. They point out possible secondary, at times even undesirable impacts of policies and innovations designed to combat the AIDS epidemic.[12] More similar to our study, Lakdawalla, Sood, and Goldman (2006) examine behavioral responses to HAART, showing that HAART increased risky sex behavior among the HIV+. They argue that this increase in risky behavior could potentially reduce the welfare of HIV− individuals, who are more likely to be matched with an HIV+ partner. We formally model this possible downside of HAART. In addition, our model is able to capture other effects of HAART, including higher continuation payoffs for HIV− agents along with lower infection rates, both of which would benefit uninfected men. As these countervailing dynamics are captured within a single framework, we can weigh their relative contribution to the total value of HAART. In particular, we can draw conclusions about whether the net effect of HAART (which incorporates possible downsides arising from increased sexual behavior of HIV+ men) is beneficial to HIV− men.

Using our estimated model, we show that the value of HAART for HIV− agents corresponds to about 7.2 years of life. We find that had HAART only improved the health of HIV+ men, but not affected seroconversion, it would be worth about 1.4 years of life. For HIV− men, we also show increases in riskier sex behaviors as a result of lower infection rates and a lower cost of HIV infection. For HIV+ men as well we see increases in sexual behavior. One reason is improvements to health. However, we also find increases in risky behavior for HIV+ men with the same level of health. We explain these increases as HIV+ men's response to the disinhibition effect of HIV− men, which lowered search costs for finding a partner willing to engage in risky sex. We also isolate the impact of market-level reactions to HAART on the value of the innovation. If HIV+ men maintained their pre-HAART sex behavior in the post-HAART era, the value of HAART would be about 4% larger for HIV− men. This relatively small incremental value reflects how the behavioral impact of HIV+ men on the welfare of HIV− men is limited once HAART is introduced. The reason is that HAART made HIV infection less costly and less likely. On the other hand, if HIV+ men engaged in risky behavior at post-HAART rates, but in the pre-HAART era when HIV infection is more likely and very costly, the resulting welfare loss for uninfected agents is

[11] Numerous studies estimate the price elasticity of demand for cigarettes and alcohol (Cawley and Ruhm, 2012), some incorporating persistent behavior through models of rational addiction (Becker and Murphy, 1988). Koszegi and Gruber (2001) demonstrate that agents exhibit forward-looking behavior in anticipation of cigarette price changes.

[12] Later contributions to this literature include: Mechoulan (2007), Auld (2003), and Auld (2006).

high: amounting to about 22% of the total value of HAART. In other words, in a scenario where HIV is still relatively likely and very costly, the impact of HIV+ men's behavior on the welfare of HIV− men through increased infection probabilities is much stronger.

Noting that approximately $600 million dollars are spent each year on the search for an effective HIV vaccine (Lamourelle et al., 2005), we also use our estimated structural model to conduct a counterfactual policy experiment to ascertain the incremental value of a perfect vaccine. A perfect vaccine essentially strengthens the HAART-induced reduction in transmission rates, driving seroconversion probability to zero. We find that for most HIV− men, a perfect vaccine is worth 1.8 years of life over and above the value generated by HAART. In other words, from the perspective of uninfected agents, the incremental value of a fully functional vaccine is comparable to the option value of HAART. Using the equilibrium framework, we also show that a vaccine generates about 9% more value than HAART from the perspective of HIV+ men. This is due to the disinhibition effect on HIV− men on the market, which makes the effort costs of finding a partner lower for HIV+ men. Notice that, by taking account of market-equilibrium effects, our model is able to capture the value of a vaccine (intended solely for HIV− men) from the perspective of HIV+ men. This is analogous to our findings on the option value: by taking account of dynamic decision-making, our model is also able to capture the value of HAART (intended for HIV+ men) from the perspective of HIV− men.

The remainder of this article is organized as follows: Section 2 introduces the data. In Section 3 we specify a dynamic discrete choice model of sexual behavior. In Section 4, we discuss the estimation and identification of our model. Section 5 describes preference parameters. In section 6, we discuss how individuals on both sides of the market best respond to the optimal behavior of agents on the other side of the market and then describe equilibrium behavior. Next, we use the estimated model to measure value of HAART and to evaluate a counterfactual, fully-effective vaccine in Section 7. Section 8 concludes.

2 Data and Preliminary Analysis

2.1 Individual-Level Data

To measure individual health trajectories, including survival, infection and HIV-induced health deterioration (AIDS), we use the public data set from the Multi-Center AIDS Cohort Study (MACS). The MACS is an ongoing longitudinal investigation (beginning in 1984) of HIV infection in men who have sex with men (MSM) conducted at four sites: Baltimore,

Chicago, Pittsburgh and Los Angeles.[13] At each semi-annual visit, survey data are collected on: sex behaviors (including condom usage), demographics and psychosocial characteristics. In addition, blood tests are administered to objectively measure health status (including CD4 count, a proxy for immune system health), serostatus (whether an agent is HIV+) and seroconversion (becoming infected). These data allow us to explicitly map seroconversion to sexual behavior. The MACS also measures uptake of HAART, which means we can observe how the diffusion of a new medical technology affects behavior.

The full MACS data set contains information on 5,622 subjects at 41 possible visits for a total of 98,886 observations (in the form of subject-visits). We exclude observations where date-of-seroconversion is unclear, attrition is non-random or if data are missing, leaving 57,276 observations and 4,028 individuals.[14] We also drop observations from the first 11 visits due to inconsistent condom usage questions and exclude the last two sample periods since deaths are often confirmed up to a year after they occur and we want to avoid confounding death with attrition. After these exclusions, we arrive at our analysis sample of 2,426 individuals over 27 half-year periods (mid-1989 until mid-2002), constituting 30,274 observations. Each individual is observed on average 12 times, ranging from 1 to 27.[15] Summary statistics are reported in Table 1.

According to Table 1, the analysis sample consists of men who were infected prior to the start of MACS data collection in 1984, which corresponds to the start of the epidemic (52% of the sample), men who remained HIV− throughout the sample period (32%) along with observed seroconverters, who are men we observe becoming infected with HIV (16%). At the first sample period (mid-1989), average age is about 30, ranging from 20 to 75. 53% of HIV+ men are college educated versus about 65% of HIV− men. About 18% are non-white, with higher numbers among infected men. These statistics point to well-established links

[13]Data in this manuscript were collected by the Multi-Center AIDS Cohort Study (MACS) with centers (Principal Investigators) at The Johns Hopkins Bloomberg School of Public Health (Joseph B. Margolick, Lisa P. Jacobson), Howard Brown Health Center, Feinberg School of Medicine, Northwestern University, and Cook County Bureau of Health Services (John P. Phair, Steven M. Wolinsky), University of California, Los Angeles (Roger Detels), and University of Pittsburgh (Charles R. Rinaldo). The MACS is funded by the National Institute of Allergy and Infectious Diseases, with additional supplemental funding from the National Cancer Institute. UO1-AI-35042, 5-MO1-RR-00052 (GCRC), UO1-AI-35043, UO1-AI-35039, UO1-AI-35040, UO1-AI-35041. Website located at http://www.statepi.jhsph.edu/macs/macs.html.

[14]A number of HIV− subjects were administratively censored between 1992-1994. The censoring occurred as follows: riskier HIV− men were kept on and less risky HIV− men were randomly censored. We circumvent problems associated with resulting potential selection bias by dropping censored individuals from the analysis for the entire sample period. A preferable strategy would be to keep censored individuals at earlier periods and control for the risky behaviors leading to censorship. However, the particular algorithm used to determine which subjects were risky enough to avoid censorship is unavailable.

[15]Since many sample subjects died between 1984 and 1989, dropping these initial periods lowers the sample size considerably. Of our analysis sample, 2,107 joined the study in 1984 and the remaining 319 joined between 1987-1989, also prior to our analysis sample period.

between health and socioeconomic status (here proxied with race and education). About 35% of agents participated in anonymous sex prior to the start of the sample period, defined as intercourse with a partner who is unknown and not able to be found. Anonymous sex captures variation in preferences for sex.

CD4 count is an objective measure of immune system health for HIV+ men. It is defined as the number of white blood cells per cubic millimeter of blood. A normal range absent HIV infection is between 500 and 1500. A count below 500 indicates that the immune system has begun to deteriorate due to HIV, but can still fight off routine infections and so the individual is not symptomatic. When CD4 count drops below 300, the immune system becomes unable to fight off routine infections, a condition known as AIDS. Henceforth, we refer to HIV+ men with CD4 counts of 300 or above as having a high CD4 count and those with counts below 300 as exhibiting a low or "AIDS-level" CD4 count. In Table 1, we see that average CD4 for HIV+ men prior to HAART is below 500, whereas for HIV− men it is about 1000. For HIV+ men, CD4 counts rose after HAART was introduced, though not for men who are observed seroconverters. This reflects how observed seroconverters, prior to HAART, had not yet been infected long enough to suffer the debilitating effects of HIV infection. The probability of surviving until the end of the sample period conditional on being HIV+ at baseline is 43%. It is higher (53%) for seroconverters who were infected later. Conditional on survival until HAART, 83% of men infected in 1984 and 83% of men who were infected later on survive until the end of the sample period. In other words, given survival until HAART, the likelihood of survival until the end of the sample period is not only much higher but is also independent of how long the individual has been infected. These patterns underscore both the severity of HIV infection and the effectiveness of HAART.

2.2 Market-Level Data

We augment the MACS data with market-level data using county-level health statistics on average HIV prevalence rates during our sample period obtained from the four cities from which the MACS sample are drawn. The reason we need additional data is that the MACS data set over-samples HIV+ individuals, which means that they do not represent an agent's knowledge of HIV prevalence in the market for partners they face. To see why this is a problem, write:

$$\begin{aligned} &\text{P[HIV Status} \wedge \text{Sex Behavior} \mid \text{MSM]} \quad = \\ &\text{P[Sex Behavior} \mid \text{HIV Status} \wedge \text{MSM]} \quad \times \quad \text{P[HIV Status} \mid \text{MSM]} \end{aligned} \tag{1}$$

The quantity on the first line represents HIV status on the MSM market, which agents use to form expectations over infection rates when information about partner status is limited. The first quantity in the second line of equation (1) can be computed using the MACS data. The second quantity, which is the probability of HIV status conditional on being MSM, is computed using county-level health statistics. A summary of market-level data used in our analysis is found in Table 2 and we plot the proportion of MSM who are in each HIV+ health status over time in Figure 1.[16]

The first column in Table 2 shows the probability in each year that a random partner on the MSM sex market is HIV+ for the four cities in our analysis. This number increases over time, which reflects higher HIV prevalence in the U.S. It also reflects how medications improved so that HIV+ people remained alive. Attenuating this rise are increases in the rates of HIV among non-MSM so that over time the proportion of new HIV cases were less likely to be among MSM.[17] In the second and third columns of Table 2, we calculate the probability that a male sex partner on the MSM market is HIV+ with a high or low CD4 count. By 2003 the vast majority of MSM HIV+ men have high CD4 counts. This is explained by the diffusion of HAART, summarized in Columns 4-6. In Column 4, we present average HAART use over time for HIV+ MSM. In Columns 5-6, we do the same men with high and low CD4 counts, respectively. After 1996, HAART use grew dramatically from about 0% to nearly 70% among HIV+ men. Not surprisingly, low-CD4 count infected men started to use HAART earlier. After a couple of years, however, the probability of using HAART is about the same for both groups.

2.3 The Impact of HAART on Health and Risky Behavior

We treat HAART introduction as a quasi-experiment.[18] Two observations justify this approach. First, HAART was not a single medication developed and improved over time such that subjects might update their beliefs and anticipate higher future efficacy. Rather, HAART introduction was abrupt and many components of HAART already existed prior to 1996. The key insight involved the union of several existing technologies, none of which was particularly effective on its own. Second, subject reports from survey questions asking about their hopefulness about the future are not consistent with anticipation of HAART.

[16] Further details on how market quantities are computed are found in Appendix A, Table A1.

[17] Indeed, virtually all HIV infections in the early years of the AIDS epidemic were among MSM. By the early 2000s, about half of all new infections were not.

[18] This assumption underscores the need for caution in applying the framework developed in this paper to cases where medical innovation is anticipated. A possible generalization of the current model would permit agents to anticipate an improvement in medical technology, similar to Keane and Wolpin (2002) and Van der Klaauw and Wolpin (2008), who study models where agents form expectations over policy changes.

Specifically, one in a battery of questions meant to assess depression asks subjects how often in the week preceding their interview they felt hopeful about the future.[19] Figure 2 plots the probability that subjects answer, "all or most of the time" over time for HIV+ men. The pre-HAART flat (or even downward) trend is followed by a break and reversal coinciding with HAART introduction. If the effectiveness of HAART had been anticipated, this upward shift in hopefulness would likely have occurred before HAART.[20]

The introduction and diffusion of HAART had two direct effects. It improved health for HIV+ individuals and it lowered infection rates, which benefitted HIV− individuals. In Table 3, we show health state transitions from one period to the next for the pre-HAART and post-HAART eras. After HAART, the likelihood of AIDS declined while the likelihood of recovering from AIDS rose. Moreover, seroconversion rates in the post-HAART era are about one quarter their pre-HAART levels. This drop can be explained by the viral suppressing effect of HAART. The magnitude of this effect is considerable. Clinical research has shown that agents using HAART are virtually non-infectious (Attia et al., 2009).[21]

HIV− men can affect their likelihood of infection through their choice of risky sexual behaviors. We consider three sexual behaviors that agents engage in for a given six month interval and report retrospectively.[22] These options are (i) non-anal sex (no intercourse or oral intercourse only), (ii) anal intercourse with a single partner or with multiple partners where condoms are used with at least one of the partners consistently and (iii) anal sex with multiple partners where condoms are not used consistently with any partner. From the perspective of HIV− agents, these behaviors are increasingly risky. These categories capture two ways that anal sex can be safer. Agents can be monogamous and not use condoms or be non-monogamous, but use condoms consistently. Not using either precaution means that the risk of HIV infection is high. Henceforth, we refer to sex behavior (i) as "low-risk sex", (ii) as "moderate-risk sex" and (iii) as "high-risk sex".[23]

[19] Questions are from a depression screening test known as the Center for Epidemiological Studies Depression (or CES-D) scale. See, for example, Ostrow et al. (1989), for an example of CES-D scale use with the MACS data set.

[20] In supplementary results available from the authors, we show a similar trend reversal after we control for health, i.e., hopefulness increases do not simply reflect HAART-induced health improvements.

[21] The study finds 5 infections for one hundred person-years among people reporting use of HAART.

[22] A key assumption we make is that data on sex behavior are accurate and that reporting accuracy did not change with HAART introduction. Agents potentially under-report or over-report riskier behavior due to stigmatization, a possibility analyzed in Gersovitz et al. (1998). In defense of our assumption, we find that "riskier" reported sex leads to a higher likelihood of infection, which was objectively measured using blood tests. Further, MACS data collection, including sensitivity training of those collecting data, was carefully designed to maximize reporting accuracy (Kaslow and Ostrow, 1987).

[23] In fact, use of condoms can misleadingly appear "risky" since men who limit their sexual activity to long-term monogamous partners use condoms less frequently than men who have many partners. We do not observe in all sample periods whether a man who does not use condoms is having sex with a monogamous

Table 4 shows sex behavior choices by health status in the pre- and post-HAART eras. Two patterns emerge. First, healthier men have more anal intercourse, where anal intercourse means individuals choose either moderate or high-risk sex. This is partly explained by health-state-dependent utility, i.e., strenuous sex is costlier when men are sick. It may also reflect higher search costs for sicker men. Second, more anal sex occurred in the post-HAART era. Table 5 shows how men switch from one behavior to another. The post-HAART rise in riskier types of sex is largely driven by higher persistence in the high-risk option. A drawback with Tables 4 and 5 is that they report unconditional probabilities that obfuscate important countervailing dynamics. Since we observe an aging cohort—and age is associated with less anal sex—these probabilities under-estimate within-age-group shifts towards moderate and high-risk sex behavior after HAART is introduced. To illustrate this point, we plot residuals from a regression of individuals engaging in anal sex over time controlling for age and health (Figure 3). The figure essentially controls for composition effects in the form of aging or health improvements. The figure shows that anal sex occurred more frequently after HAART for men with similar ages and in the same health categories. Further evidence of this shift is found in Table 6, where we present estimates from a multinomial logistic regression of sex behavior onto health, age and other socio-demographic variables. Estimates show that the post-HAART shift towards moderate and high-risk sex is significant.[24]

Having established that both HIV+ and HIV− agents engage in riskier sex after HAART, we now ask why. In Table 7, we regress HIV− men's sex behavior onto the percentage of HIV+ agents using HAART in each time period (plotted in Figure 4). Estimates confirm that HIV− men engaged in more risky sex as HAART use rose. There are two reasons for this. First, seroconversion declined since HAART renders HIV+ agents effectively non-infectious. In Table 8, we present estimates of seroconversion regressed onto sexual behavior and HAART use of HIV+ men. Notice, if we only regress onto the proportion of HIV+ agents having anal sex (Column [2]), the coefficient is negative since we capture how HAART use lowered seroconversion rates, but also led to higher rates of riskier sex of HIV+ men. If we only regress onto HAART use, the coefficient is negative (Column [3]). Further, if we control for HAART use, the coefficient on the portion of HIV+ agents having anal sex becomes positive, though insignificant (Column [4]). These results underscore how HIV+ agents' behavior (both HAART use and risky sex) affects HIV infection rates. A second reason why HIV− agents shift towards risky sex is that the long-run cost of HIV infection is lower after

partner (versus having unprotected sex with an anonymous partner). Therefore, we reserve the distinction of "high-risk sex" for the choice to have multiple partners and to not using condoms consistently with any one of them.

[24]Qualitative results are robust to the inclusion of fixed effects if we run a binary choice logit model where options are anal sex versus none or oral sex only. Results are available from the authors.

HAART. This can be seen in Figures 5 and 6, which plot survival and health for infected men. Notice that the likelihood of dying for HIV+ men declines precipitously after HAART is introduced, as does the probability of having (and failing to recover from) AIDS. In other words, being HIV+ was no longer as terrible as it had been prior to HAART and HIV− men responded by choosing riskier sex more often.

Finally, we ask why HIV+ men also increased their risky sexual behavior after HAART. One reason, examined in Dilley, Woods, and McFarland (1997) and in Lakdawalla, Sood, and Goldman (2006), is that their health improved. However, this explanation can only be part of the story. In Table 6, we show that post-HAART increases in risky sex among HIV+ men occur *within* each health category. A possible concern is that changes in HIV+ men's sex behavior attributed to HAART are instead due to rises in CD4 count within each health category (e.g., men who have relatively high CD4 counts, but are still below the AIDS cutoff). To control for this possibility, in the multinomial logistic regressions presented in Table 6, we have included health-state-specific continuous measures of CD4 count. There is some evidence of within-health effects for men with AIDS. However, despite the inclusion of these additional regressors, all post-HAART dummy variables interacted with health categories remain positive and significant.[25] We argue that, as HIV− men substitute towards riskier sex (and information about partners is limited), HIV+ men face lower effort costs of finding willing partners for risky sex. To capture this effort cost in the model we estimate, we specify static payoffs for sex behavior as a function of the optimal behavior of HIV− men. We now turn to the specification of the theoretical model to be estimated.

3 Model

In this section, we specify the dynamic choice model of risky sex behavior in the context of HIV and the diffusion of a new medical technology, HAART. Our goal is to estimate parameters in a model that is capable of quantifying the various effects of HAART discussed in the previous section. In Section 3.1, we specify the choice set and state variables. State variables include time-specific market-level quantities that affect agent choices and outcomes, such as the proportion of homosexual men who are HIV+ and the proportion of HIV+ men engaged in high-risk sex. In Section 3.2, we introduce the value function and specify its

[25] In additional specifications, available upon request, we also include interactions between health states and the dummy variable for having a college degree to control for the possibility that increased sex behavior is driven solely by college educated men, who might have better access to new medical technology. With the inclusion of these controls, the post-HAART dummy variable coefficients remain significant and positive. In other words, there is clear evidence that HIV+ men increase their sex behavior after HAART for reasons that are not attributable solely to health improvements.

components. Once the dynamic program has been fully specified, we describe how it is solved in Section 3.3. Description of the solution includes defining the following equilibrium constraint: agent beliefs about market-level behavior accord with optimal choices implied by the model. Estimation of the model is discussed in Section 4.

3.1 Choices and States

In this section, we specify the choice set along with the vector of state variables, which includes market-level quantities that affect agent behavior. At time t, agent i chooses a sex behavior $d_{it} \in \mathcal{D} \equiv \{1, \ldots, D\}$. Sex behavior options are:

$$d_{it} = \begin{cases} 0 & \text{Low-risk sex} \\ 1 & \text{Moderate-risk sex} \\ 2 & \text{High-risk sex} \end{cases} \tag{2}$$

Recall from the previous section that "low-risk sex" means that agents have oral sex only or no intercourse. "Moderate-risk sex" means that agents have anal sex, but take at least one precaution, including limiting themselves to one partner or consistently using condoms with at least one partner. "High-risk sex" means anal sex where neither precaution is taken. Next, we define a vector of observed state variables for agent i at time t and denote it Z_{it}. The set of observed state variables contained in Z_{it} includes:

$$\begin{aligned} X_{it} \in \{X^1, \ldots, X^M\} &: \text{Non-health state variables at } t \\ h_{it} \in \{0, 1, 2\} &: \text{Health status at } t \\ H_{it} \in \{0, 1\} &: \text{HAART use at time } t \\ B_{it} \in \{0, 1\} &: \text{Death before the start of } t \\ \bar{h}_t &: \text{Average health status shares at } t \\ \bar{H}_t &: \text{Average HAART use shares for HIV+ agents at } t \end{aligned} \tag{3}$$

Here, M is the number of possible combinations of non-health state variables and is finite. Non-health state variables include age ($A_{it} \in \{25, 30, \ldots, 75\}$), an indicator for a college degree or higher (C_i), an indicator for having reported anonymous sex prior to the start of the sample period (N_i) along with a dummy variable taking the value one for non-white

individuals (R_i).[26] Further, agents at any point in time are in one of three health states, h_{it}

$$h_{it} = \begin{cases} 0 & \text{Uninfected (HIV}-\text{)} \\ 1 & \text{Infected (HIV+) with a healthy (high) CD4 count} \\ 2 & \text{Infected (HIV+) with an unhealthy (low, AIDS-level) CD4 count} \end{cases} \tag{4}$$

Here, $h_{it} = 2$ means CD4 count is low enough so that the immune system is less able to fight off what would normally be routine infections, a condition known as AIDS. Agent decisions are also a function of market-level quantities that describe the evolution of HIV infection rates along with the diffusion of HAART, represented by the final two variables in the state space. $\bar{h}_t$ is a (1×3) vector where the first, second and third elements are proportions of the population of MSM in health states 0, 1 and 2, respectively. $\bar{H}_t$ is a (1×2) vector of HAART use where the first and second elements are the proportion of HIV+ in health states 1 and 2, respectively, who are using HAART. Finally, B_{it} indicates dying before the start of period t. In summary, the vector of observed state variables for individual i at time t is defined as follows:

$$Z_{it} = \left\langle X_{it}, h_{it}, H_{it}, B_{it}, \bar{h}_t, \bar{H}_t \right\rangle . \tag{5}$$

We also define a vector of unobserved (to researchers) state variables that are choice specific:

$$\epsilon_{it} = \left\langle \epsilon_{i1t}, \ldots, \epsilon_{iDt} \right\rangle . \tag{6}$$

In the model, agent choices and outcomes are influenced by the sex behavior decisions of other agents. For example, if HIV+ agents on average shift towards a particular sex behavior, an HIV− man engaging in the same sex behavior is more likely to match with an HIV+ agent and become infected. To summarize average behaviors at the market-level, we define $\bar{P}_t$, which is a (3×3) matrix, where element (h, j) of $\bar{P}_t$ is the proportion of agents in health state h who choose sexual behavior $d_{it} = j$ in period t and is denoted $\mathrm{P}^j(h, \bar{h}_t, \bar{H}_t)$. This market-level probability is a weighted average over X_{it} so that:

$$\mathrm{P}^j(h, \bar{h}_t, \bar{H}_t) = \sum_{X_{it}} \mathrm{P}_t(X_{it}|h) \times \mathrm{P}^j(X_{it}, h, \bar{h}_t, \bar{H}_t), \tag{7}$$

where $\mathrm{P}_t(X_{it}|h)$ is the health-state conditional distribution of observed state variables in period t. Since older men tend to engage in less anal sex as they age, we also compute average behavior in each health state separately for men who are younger than 40 ($\bar{P}_t^{Younger}$)

[26] Categorizing age in increments of 5 years decreases the size of the state space. For estimation, one-period-ahead value is a weighted average of value given aging five years and value given remaining in the same age category.

versus older than 40 ($\bar{P}_t^{Older}$) and collect the resulting age-specific average choice probabilities into a matrix denoted $\bar{P}_t^{Age}$.

3.2 The Value Function

In this section, we introduce the choice-specific value function, which is a function of current-period flow utility along with expected future utility conditional on current-period states and decisions. We start by writing the value function in a fairly general form and then specify each component in detail, assuming that $B_{it} = 0$, i.e., the agent is alive at the start of period t. For state variables Z, the value function for choice $d = j$ is:

$$\begin{aligned} V_j(Z) &= u_j(h, X) - e_j\left(h, \bar{h}, \bar{P}^{Age}\right) + \epsilon_j \\ &+ \textstyle\sum_{Z'} f_j\left(Z'|Z, \bar{P}\right) \times \beta \Big[\mathrm{E}_\epsilon \left[\max_k \left\{V_k\left(Z'\right) + \epsilon'_k\right\}\right] \times (1 - B') \\ &+ \nu \times B' \Big]. \end{aligned} \tag{8}$$

To save on notation, we have suppressed individual and time indices, do not condition on parameters to be estimated and mark one-period-ahead values of state variables with a prime. The first line on the right-hand-side of equation (8) contains expressions for current-period returns for each choice j. This includes the flow utility function $u_j(\cdot)$, which captures preferences for sex, is choice-specific and is a function of the agent's health h and non-health state variables X. Current-period returns also include the effort function $e_j(\cdot)$, capturing how agents potentially face variation in the cost of finding partners, which is a function of average sex behavior choices for each age group $\bar{P}^{Age}$ interacted with health state market shares $\bar{h}$.

The second line of equation (8) contains a discounted weighted average of continuation payoffs. The first component is the probability distribution function over future state variables and market-level quantities, denoted $f_j(\cdot|\cdot)$, which is choice-specific and also a function of sex choices of other agents at the market level $\bar{P}$. This captures how HIV− agents face a probability of infection that is a function of their own sex behavior j along with the average sex behavior of other agents $\bar{P}$. With regard to the probability distributions over future states and over other agents' behavior, we assume here, and throughout this study, that subjective probabilities (beliefs) are consistent with objective probabilities. This assumption amounts to imposing rational expectations and is necessary for identification. In the Conclusion, when discussing extensions to this study, we suggest what additional data would be needed to relax rational expectations.

The second component in Line 2 of equation (8) denotes continuation payoffs, which take the form of an expected maximum since agents make optimal choices at all future states. The expectation is taken over future values of the unobserved state variable ϵ'. The agent obtains the continuation payoff only if he survives until the next period. The expression in Line 2 is therefore weighted by one minus an indicator for surviving until the next period $(1 - B')$. If he does not survive, the agent instead obtains a one time payment ν at the end of life (Line 3). ν is similar to a "scrap value" of life and is a parameter to be estimated. We now describe each component of equation (8) in greater detail.

3.2.1 Current Period Payoffs

The flow utility function for sex behavior j given health h and characteristics X is specified as:

$$\begin{aligned} u_j(h, X) &= \theta^u_{jh1} + \theta^u_{jh2} \times A \\ &+ \theta^u_{jC} C + \theta^u_{jN} N + \theta^u_{jR} R, \end{aligned} \tag{9}$$

From the first line on the right hand side, flow utility for each choice varies by health and an age-health interaction term means that age-dependent utility for each choice can vary by health state. The second line includes indicators for college degree (or more), anonymous sex prior to the sample period and race (non-white).[27] For identification, we maintain the assumption that $u_0(h, X) = 0$ so that the flow utility of no sex is zero for all health states.

The effort function captures the effort cost of finding a partner willing to engage in anal sex. Despite high rates of HIV infection among men who have sex with men, the overwhelming majority are HIV−. We focus on how HIV− men's sex behaviors affect the effort of finding willing partners for HIV+ men. For $h \in \{1, 2\}$ and for $j \in \{1, 2\}$, the effort function is:

$$e_j\left(h, \bar{h}, \bar{P}^{Age}\right) = \theta^e_{jh} \times \mathrm{P}(h = 0) \times \left[0.5 \times \bar{P}^{Younger} + 0.5 \times \bar{P}^{Older}\right] \tag{10}$$

Here, $\mathrm{P}(h = 0)$ is the first element of $\bar{h}$, representing the proportion of MSM who are HIV− in each period. This is multiplied by an expression that averages over the sex choice probabilities for younger ($\bar{P}^{Younger}$) and older ($\bar{P}^{Older}$) agents in the market to capture changes over time in the proportion of HIV− men who choose low-risk sex, taking account of how our sample is an aging cohort and men tend to engage in less anal sex as they age.[28] To fix ideas with an example, consider the effort cost of finding a partner for an HIV+ man with a high CD4

[27] The anonymous sex variable is used as a reduced-form way to capture heterogeneity in preferences for different sex behaviors.

[28] The average is weighted by 0.5 as roughly half our sample is below 40.

count ($h = 1$) choosing moderate-risk sex ($d_{it} = 1$). If θ^e_{11} is estimated to be a positive number, then the effort cost increases when more HIV− men eschew anal sex and instead choose safe sex. Positive effort is then subtracted from flow utility so that the static payoff is lower. Finally, we make the identifying assumption that effort is equal to zero for HIV− men: $e_j\left(h, \bar{h}, \bar{P}^{Age}\right) = 0$ if $h = 0$. We also normalize the effort cost for $j = 0$ to zero. Therefore, the effort function captures the effect of HIV− men's sex behavior on effort costs facing HIV+ men for moderate or high-risk sex relative to the effect on effort costs facing HIV− men.[29]

Finally, realizations of the choice-specific unobserved disturbance $\epsilon_j : \mathcal{D} \to \mathbb{R}$ are assumed to be independent across time, agents and choices and Extreme Value Type I distributed.

3.2.2 Health State Transitions

In this section, we discuss health state transitions for HIV+ and HIV− men. Once agents are HIV+, their health can fluctuate between high and low CD4 counts. These fluctuations are a function of the availability and diffusion of medical technology—in particular, the introduction and use of HAART. Health state transitions for HIV+ agents are given by

$$\mathrm{P}(h'|h, H, X), \text{where } h, h' \in \{1, 2\} \tag{11}$$

According to equation (11), transitions among health states for HIV+ agents are conditional on HAART usage. It is crucial to note that we do not explicitly model how agents optimally choose to use HAART. Instead, we use reduced-form estimates of choice probabilities taken from our sample, which we treat as optimal choice probabilities conditional on health and non-health state variables. These are denoted:

$$\mathrm{P}(H|h, X), \text{where } h, h' \in \{1, 2\} \tag{12}$$

Details of the specification are found in Section B of the Appendix. This approach places some limits on the sorts of counterfactual policy simulations that can be reasonably conducted using the estimated model, a point to which we will return in Section 6.[30] HIV−

[29] The effort costs of finding partners willing to engage in the same sex behavior for HIV− men are absorbed in the flow utility function in equation (9).

[30] It would be difficult to use the model to evaluate policies that affect agents through changes in the optimal use of HAART unless we know (or explicitly specify) how the policy would affect HAART usage. As the focus of our counterfactual policy experiments is the impact of HAART introduction and HAART use on uninfected individuals, we are less concerned with how, precisely, shifts in the use of HAART arise. Nonetheless, how infected agents choose HAART is of interest in its own right, for example, in understanding tradeoffs endemic to chronic illness. Using the MACS data set, but restricting attention to HIV+ individuals,

agents also use equation (12) to evaluate the continuation payoffs for being infected. They assume that if they become infected they will use HAART the same way that currently infected agents do and therefore face the same probability of being in a high-CD4 versus low-CD4 health state.[31]

Next, we consider health state transition probabilities for HIV− agents, i.e., the probability of seroconversion conditional on sex behavior d. We start by writing this transition as:

$$\mathrm{P}(h' = 1 | h = 0, X, d) \tag{13}$$

Written as such, seroconversion is simply a function of an individual's own sex behavior. This ignores how HIV− agents will only be infected if their partners are HIV+ and not using HAART. The chance that a partner is HIV+ depends on two factors. First, it depends on the market-level distribution of health status. Second, it depends on the likelihood of matching with a partner who is HIV+ and who also agrees to the same sexual behavior. To incorporate these factors, we rewrite equation (13) as:

$$\begin{array}{rcl} \mathrm{P}(h' = 1 | h = 0, X, d) & = & \sum_{\tilde{h}} \mathrm{P}(h' = 1 | h = 0, X, d, \tilde{h}, \tilde{H} = 0) \\ & \times & \mathrm{P}(\tilde{d} = d | \tilde{h}) \mathrm{P}(\tilde{H} = 0 | \tilde{h}) \mathrm{P}(\tilde{h}) \text{ where } \tilde{h} \in \{1, 2\}. \end{array} \tag{14}$$

Here, $\tilde{d}$ is the partner's sex behavior decision, $\tilde{H}$ is the partner's HAART usage and $\tilde{h}$ is his partner's health status. To rewrite equation (13), we have integrated out the HIV status and HAART usage of the sexual partner, given that the partner chooses the same sexual behavior.[32] This captures the following idea: when an agent chooses a sexual behavior d, he draws inferences about his partner's HIV status and HAART usage from his partner's decision to also engage in sex behavior d. To derive equation (14), we assume that HIV+ agents are nearly non-infectious when they are using HAART ($\tilde{H} = 1$), and restrict our attention to partners who are HIV+ ($\tilde{h} \in \{1, 2\}$). The joint probability of partner behavior and health is decomposed into three components, which are written on the second line of equation (14). The first component is the probability that a partner of health state $\tilde{h}$ chooses

Papageorge (2015) estimates a model of joint medication and labor supply choices that takes explicit account of the tradeoff between drug effectiveness and side effects and how this tradeoff interacts with employment decisions. The model is used to study how chronically ill individuals react to counterfactual labor market policies and medical technologies, in part through shifts in their optimal use of HAART.

[31] We attempted to relax this assumption by allowing for unobserved heterogeneity in HAART use, which would account for the possibility of dynamic selection in the subset of the population that remains uninfected with HIV. We find no evidence of this type of dynamic selection. We return to this point in section 5, where we discuss parameter estimates.

[32] There is a slight abuse of notation when we write $\mathrm{P}(\tilde{d} = d, \tilde{h}, \tilde{H})$. The reason is that a partner must choose a sexual behavior $\tilde{d}$ that accommodates d. For example, if an individual chooses anal sex (moderate or high-risk sex), we impose that he cannot match with a partner choosing oral sex only (low-risk sex).

sex behavior $\tilde{d} = d$.[33] The second component is the probability that the partner is not using HAART given his HIV status. The third component is the unconditional probability that an agent in the market is in health state $\tilde{h}$.

In equation (14), seroconversion conditional on sexual behavior d is an explicit function of partner HIV status and HAART usage that is related to partner sexual behaviors.[34] The first expression on the right hand side of equation (14) can be interpreted as the probability that an agent becomes infected, conditional on his reported sex behavior along with the health and HAART usage of his partner. It captures a type of biological rate of infection for negative agents with characteristics X, but cannot be interpreted directly as such since it averages over unobserved information, including the frequency of sexual acts, total number of partners (each, presumably, with different characteristics) and other important information that would change the likelihood of infection. The expression is therefore treated as an unknown quantity to be estimated that maps observed market factors to observed seroconversion rates. Using this idea, we can rewrite equation (14) as follows:

$$\begin{aligned} \mathrm{P}(h' = 1 | h = 0, X, d) &= \Phi^d(\tilde{h} \neq 0, \tilde{H} = 0) \\ &\times \mathrm{P}(\tilde{d} = d | \tilde{h} \neq 0) \times \mathrm{P}(\tilde{H} = 0 | \tilde{h} \neq 0) \times \mathrm{P}(\tilde{h} \neq 0). \end{aligned} \tag{15}$$

Equation (15) emphasizes the idea that Φ is an unknown parameter. Treating biological infection rates as such, however, means that we must be take great care in choosing our counterfactual policy simulations and only conduct experiments where we can be sure that Φ does not change or, alternatively, where we know precisely how Φ will shift under the counterfactual regime. One example is to conduct a policy simulation to assess the value to HIV$-$ agents of a universal vaccine that is fully effective. In that case, we know that $\Phi = 0$. We also note that, although it is in principle possible to consider both HIV+ health states separately, there is not enough independent variation between the two groups to credibly identify separate parameters Φ for each group. Therefore, we assume the same Φ for $\tilde{h} = 1$ and $\tilde{h} = 2$.[35]

We use equation (15) to connect the market-level quantities to observed rates of infection. We impose three constraints with regard to information about unobserved partners. In particular, we replace the three components on the second line of equation (15), each of

[33] Here, we have assumed that, conditional on health, use of HAART does not affect sexual behavior, a simplification that is supported by reduced form analysis.

[34] We note that integrating out partner behaviors and states as we have in equation (14) assumes that agents do not engage in assortative matching. That is, they do not sort based on their $\tilde{H}$ or on observables such as age and race.

[35] In the empirical implementation discussed in detail in Appendix C, we estimate a slope and an intercept term relating market quantities to seroconversion rates.

which contains a measure of a partners' characteristics, with expressions that do not require information about unobserved partners. The following three constraints are assumed to hold:

1. $\mathrm{P}(\tilde{d} = d|\tilde{h}) = \mathrm{P}(d|h)$, where $\mathrm{P}(d|h)$ is a component of $\bar{P}$
2. $\mathrm{P}(\tilde{H} = 0|\tilde{h}) = \mathrm{P}(H = 0|h)$
3. $\mathrm{P}(\tilde{h}) = \mathrm{P}(h)$

The first constraint imposes that choice probabilities conditional on health for unobserved partners, as will be explained in Section 3.3, be the same as the model prediction of the optimal choice of individuals under market equilibrium. The second constraint assumes that, conditional on health status, average non-use of HAART is the same for agents we observe in our MACS sample as it is for their potential sex partners. A model predicting HAART use estimated on MACS data is used to measure $\mathrm{P}(H = 0|h)$.[36] Notice that constraints 1 and 2 amount to assuming that, given the same h and X, the choices of d and H of those agents not in our sample are the same as for those in our sample. That is, there is no sample selection issue related to unobserved factors affecting sex preferences and HAART usage. The latter assumption seems very reasonable given that the sample was chosen before HAART was introduced. The third constraint imposes that the distribution of market-level health state shares be the same for individuals we observe and for their unobserved partners. As the MACS over-samples HIV+ agents, we use HIV prevalence constructed from our market-level data.

3.2.3 Remaining State Variables

In equation 8, to calculate the value of a choice, agents also form expectations over survival, modeled as a function of health, age and other controls.[37] The expectation is assumed to be consistent with the conditional survival rate that is estimated using MACS data. Parameters of the survival and HAART use functions are collected into a vector denoted θ^X. Details of the specifications are found in Section B of the Appendix. Agents also form expectations on one-period ahead average HAART use ($\bar{H}'$) and health market weights ($\bar{h}'$). As HAART introduction was not anticipated, its diffusion was rapid and the proportion of users leveled off fairly quickly and remained constant throughout the remainder of the sample period, we

[36] In particular, we start with the predicted probability of HAART usage conditional on X and h (using the reduced-form choice model defined in equation (12). Details of the specification are found in Appendix B. Then, we integrate over the joint probability of X and h to obtain the probability of HAART usage conditional on health state h.

[37] The model does not require infection with HIV for agents to die. Survival of HIV− men varies with age.

assume that before 1996 $\bar{H} = 0$. After 1996, we assume that $\bar{H}' = \bar{H}$. We also approximate beliefs over the evolution of market weights as $\bar{h}' = \bar{h}$.[38]

3.3 Choice Probabilities under Market Equilibrium

In this section, we first describe how we obtain equilibrium choice probabilities from the value function. Next, we discuss an equilibrium constraint imposed on the solution to the dynamic optimization problem. Given our assumption that the unobserved state variables ϵ_j are extreme-value type I distributed, we can write the net-of-errors choice-specific value function (which is equation (8) minus ϵ_j) as follows:

$$\begin{aligned} \bar{V}_j(Z) &= u_j(h, X) + e_j\left(\bar{h}, \bar{P}\right) \\ &+ \textstyle\sum_{Z'} f_j\left(Z'|Z, \bar{P}\right) \times \beta \Big[\left(\gamma + \log\left[\sum_k \exp\left\{V_k\left(Z'\right)\right\}\right]\right) \times (1 - B') \\ &+ \nu \times B' \Big] \end{aligned} \tag{16}$$

where γ is Euler's constant. Equation (16) is therefore a simultaneous equations system. If $\beta < 1$, it defines a contraction mapping, which permits approximation of $\bar{V}_j(\cdot)$ for each set of state variables via value function iteration. Once we have solved for $\bar{V}_j(\cdot)$, we use our assumptions on the error term ϵ_j once again to construct choice probabilities conditional on state variables. In particular, we define optimal choice probabilities as follows:

$$\phi_j(Z) \equiv \mathrm{P}(d = j|Z) = \frac{\exp\left\{\bar{V}_j\left(Z\right)\right\}}{\sum_{k=1}^{D} \exp\left\{\bar{V}_k\left(Z\right)\right\}} \tag{17}$$

Equation (17) defines choice probabilities as a function of value functions minus the choice-specific disturbances ϵ_j. The estimation procedure outlined in Section 4 will include matching these theoretical choice probabilities to analogous choice probabilities from our sample.

Next, we impose the following equilibrium constraint.

$$\mathrm{P}^j(X_{it}, h_{it}, \bar{h}_t, \bar{H}_t) = \phi_j(Z) \tag{18}$$

Recall that $\mathrm{P}^j(X_{it}, h_{it}, \bar{h}_t, \bar{H}_t)$ are choice probabilities that are used in equation (7) to construct the market-level average choice probabilities $\bar{P}_t$ and $\bar{P}_t^{Age}$, which enter into the agent's dynamic choice problem through the effort function for HIV+ agents (equation (10)) and the expected seroconversion probability function for HIV− agents (equation (15)). The equi-

[38] In other words, agents assume that one-period-ahead HAART use and proportions of MSM in each health state are the same as their current-period values.

librium constraint means that the market-level choice probabilities that factor into agent decision-making are the same choice probabilities implied by the optimal choice model. Having imposed the constraints as detailed above, we rewrite equation (17) conditioning on parameters of the model to be estimated, which we collect into a vector denoted Ξ:

$$\Xi \equiv \left\langle \theta^u, \theta^e, \theta^X \right\rangle \tag{19}$$

Optimal choice probabilities when the market is at equilibrium are therefore specified as:

$$\phi_j(Z;\Xi) = \frac{\exp\left\{\bar{V}_j\left(Z, \phi(Z;\Xi)\right)\right\}}{\sum_{k=1}^{D} \exp\left\{\bar{V}_k\left(Z, \phi(Z;\Xi)\right)\right\}}. \tag{20}$$

Equation (20) captures the idea that in the context of infectious disease, medical innovations will change individual behaviors and thus affect the infection rate. Anticipating these changes, individuals will further alter their behaviors, leading to further changes at the market level. The equilibrium is represented by the fixed points $\phi(Z;\Xi)$ in equation (20). It defines a simultaneous equation system since $\phi(\cdot)$ is on both sides of the equation, which will affect the estimation algorithm. Also notice that the vector Ξ is on both sides of equation (20), which reflects how structural parameters not only govern optimal choices, but also govern how agents optimally respond to equilibrium behavior.

4 Estimation and Identification

In this section, we briefly summarize how the vector of parameters Ξ is estimated via maximum likelihood and relegate further details on the estimation algorithm to Appendix C. Thereafter, we discuss how model parameters are identified using the data described in Section 2. In describing estimation, we note that some quantities used in the algorithm are directly computed in a separate first step since they do not change as model parameters change. These include market weights from county-level health statistics $\bar{h}_t$ along with HAART choice probabilities $\bar{H}_t$ and transition probabilities for HIV+ agents among health states from the MACS data set.

In the second step, we use these fixed quantities in estimating model parameters in an iterative way. First, we impose an initial guess of time-specific, market-level choice probabilities for HIV− and HIV+ men, $\bar{P}_t$ and $\bar{P}_t^{Age}$. As discussed in the model section, HIV− men's choices affect the effort function for HIV+ men, whereas HIV+ men's choices affect infection rates. Given a trial value of Ξ, along with our initial guess of aggregate choice probabilities, we construct static payoffs, including utility $u_j(\cdot)$ and effort $e_j(\cdot)$, for each

choice and combination of state variables. Next, we regress seroconversion onto market level quantities (HAART use, time-specific choice probabilities and market weights) in equation (15), which yields estimates for Φ. With these estimates, along with HIV+ health state transition probabilities computed in the first step, we can obtain choice-specific, net-of-error value functions for each state $\bar{V}_j(\cdot)$ from equation (16). Finally, we use assumptions on the disturbances ϵ_j to obtain choice probabilities from these value functions, ϕ_j in equation (17).

Choice probabilities ϕ_j reflect optimal behavior given model parameters along with our initial guess of $\bar{P}_t$ and $\bar{P}_t^{Age}$. However, ϕ_j when aggregated to the market-level may be different from initial guesses used in modeling agent expectations. In that case, we have calculated optimal behavior in light of beliefs about equilibrium behavior that are not consistent with equilibrium behavior implied by the model ϕ_j. If so, we aggregate over optimal choice probabilities implied by the model, ϕ_j, for each period and then use these to update our initial guess for $\bar{P}_t$ and $\bar{P}_t^{Age}$ and repeat the second step of the estimation algorithm. Iteration continues until the following condition is met: in every time period, optimal behavior implied by the model ϕ_j is consistent with agent expectations about optimal behavior $\bar{P}_t$ and $\bar{P}_t^{Age}$. This gives us the equilibrium ϕ_j conditional on the trial value of Ξ. At the "outer" loop, we repeat the above procedures to compute equilibrium choice probabilities until the likelihood function is maximized.[39]

4.1 Identification

Here, we discuss how variation in the data is used to identify model parameters. We begin with HIV+ men. For them, sex behavior does not affect continuation payoffs. Therefore, after normalizing the flow utility of low-risk sex to zero, we can identify flow utility parameters of HIV+ men from their observed health-state-dependent sex behavior choices. Using state-to-state transition probabilities identified from analogous moments in the data, we can compute lifetime utility for HIV+ men.

For HIV− men, the challenge is to distinguish between flow utility and continuation payoffs. We cannot do so solely from observed choices since sex behavior affects both. In particular, the value function for HIV− men is the utility of current sex behaviors plus the continuation payoffs from being uninfected and infected, weighted by the one-period-ahead transition probabilities as functions of current sex behaviors. To capture continuation payoffs, we exploit the rational expectations assumption and the market equilibrium constraint, both of which allow us to calculate the one-period-ahead probabilities. Conditional on their becoming infected in the next period, we can compute HIV− men's lifetime utility using

[39] How we compute standard errors is also discussed in Appendix C.

HIV+ men's flow utility parameters, which were identified from HIV+ men's choices. Once we have controlled for the one-period-ahead probabilities and the continuation payoffs of being HIV+ as of the next period, we can use HIV− men's observed sex behavior choices to identify their flow utility parameters and the continuation payoff of being HIV− in the next period.

HAART also led to shifts in equilibrium behavior. Under the identifying assumption that state-dependent preferences (θ^U) are stable across time, HAART-induced increases in risky sex behaviors among HIV+ men can be explained by health improvements of HIV+ men and the ease of finding partners engaged in the same sex behavior. To separately estimate these effects, we specify an effort function, which is a mapping from HIV− men's behavior to HIV+ men's payoffs for anal sex. The changes in average equilibrium behaviors of the HIV− men after HAART help to identify the parameters of the effort function. The observed health improvements of HIV+ men after HAART was introduced help to identify the health-specific utility parameters for sex behaviors among HIV+ men.

We exploit shifts in the behavior of HIV− men before and after HAART to identify the scrap value of life. After HAART, the probabilities of infection and of death once infected are reduced. Changes to observed choice probabilities of moderate and high-risk sex that cannot be explained by these observed changes in infection rates are attributable to reduced death probability and help to identify the scrap value of life. In other words, the scrap value of life is identified through HAART-induced changes in behavior of HIV− men, which cannot be explained by changes in continuation payoffs from a lower likelihood and utility cost of HIV.

One potential issue related to using the contraction mapping in equation (20) is that there may be multiple fixed points, i.e. multiple ϕ's that support the equation for a given set of parameters. Though we cannot rule out multiplicity of equilibria given the structure of the model, our estimation algorithm chooses the set of choice probabilities that best fits the data and also satisfies the equilibrium constraint. Our results show, however, that multiple equilibria do not exist under the estimated parameters. This will be helpful since we do not have to be concerned about which equilibrium will arise under counterfactual policy experiments. Further details are in the next section.

5 Parameter Estimates and Sex Behavior Dynamics

In this section, we discuss parameter estimates and then use the estimated model to simulate sex behavior to assess fit and to simulate sex behavior over time. We begin with parameters

from the first stage of estimation, including those governing the decision to go onto HAART for HIV+ men with high and low CD4 counts (Table 9) transitions among health states for HIV+ men (Tables 10 and 11) and survival (Table 12). The first stage estimates are largely consistent with the reduced-form, preliminary estimates discussed in Section 2. Use of HAART is a function of the diffusion of the new technology, captured by time dummy variables, along with observable socio-demographic variables. Recall, HIV− men form beliefs over the choices they would make if they were HIV+. For any given time period, we assume that HIV− men expect that they would use HAART at the same average rate as observably similar HIV+ agents.[40] The time dummy variables in Table 9 show that the diffusion of HAART was rapid for both high and low CD4 count men, occurring largely between 1995 and 1997 and stabilizing thereafter. This is consistent with the patten depicted in Figure 4. Further, within both health categories, older, college educated and white men are more likely to use HAART, which may reflect better access to medical technology and a higher willingness to invest in health among men with higher expected earnings. Men who have reported anonymous sex are less likely to use HAART, though this relationship is only significant for men with AIDS.

Parameters governing how HIV+ men transition between high and low CD4 counts are found in Table 10 (high to low CD4) and Table 11 (low to high CD4). According to estimates, older men are more likely to transition in either direction as are men who previously engaged in anonymous sex. College educated men are less likely to transition to AIDS and are more likely to recover. This link is in line with previous research linking socio-economic status and health. Finally, HAART use not only prevents AIDS, but also induces recovery from AIDS. The importance of recovering from AIDS-level CD4 counts is highlighted in Table 12, where we present estimates on survival. According to our estimates, men who are HIV+ with a high CD4 count are about as likely to die as men who are HIV−. Men who are HIV+ with a low (AIDS-level) CD4 count are far less likely to survive until the next period. Further, age is associated with a lower survival probability as is lack of a college degree, the latter reflecting once again a strong relationship between socioeconomic status and health. Finally, we note that once we have controlled for health status and education, race has no effect on survival probability.

[40] We also estimated a model permitting permanent unobserved heterogeneity in HAART usage to capture how men who are HIV− might, due to dynamic selection, systematically differ in their expected use of HAART versus men who are currently infected. In particular, we allowed two HAART usage types, but found that estimated parameters placed virtually all weight on the unobserved type using HAART according to estimated averages for HIV+ men. This finding is broadly consistent with previous work showing little evidence of permanent unobserved differences in HAART use. Rather, Papageorge (2015) shows that HIV+ agents cycle on and off of HAART and this cycling is explained by observed health states, which we control for.

Next, we report results from our structural model estimation. Estimates of the parameters from a logit model linking sex behavior of HIV− men to HIV infection (equation (15)) are found in Table 13. The large and negative constants imply that seroconversion is a low-probability event. They also imply that the low-risk option is less likely than anal sex to lead to infection and that high-risk sex is more likely to result in infection than moderate-risk anal sex. The estimated slope parameters explicitly link market level quantities, including the optimal behavior of HIV+ men, to infection. The slopes are positive, implying that lower HAART usage, higher rates of either type of anal sex among HIV+ men and higher rates of HIV among MSM all lead to higher rates of seroconversion for each sex behavior. According to the estimates, the effect of these market level quantities on infection are largest for low-risk sex. However, differences in the slope parameters are not large enough to overcome differences in the constants at observed market-level quantities. The reason is that market-level quantities are all probabilities, which, when multiplied together, yield very small numbers. For example, if HAART use is 60%, HIV infection rates are about 10% and the choice of anal sex is 70% among HIV+ men, then the slope parameters are multiplied by 0.042 (compared to 1 for the constants). The effects of market-level quantities on seroconversion for sex behavior, though important, are therefore not large enough to shuffle the ordering by risk of seroconversion of the three sex behaviors.

Structural parameters from the flow utility function (equation (9)) are found in Table 14. The utility of anal sex (moderate or high-risk sex) relative to oral sex only or no sex (low-risk sex) declines when agent health declines or when agents age, reflecting how physical fitness influences enjoyment of sex. College educated men and white men prefer moderate-risk sex (compared to low-risk sex) more strongly than their less educated or non-white counterparts. Race and education do not have a significant systematic impact on preferences over high-risk sex. However, previous engagement in anonymous sex reflects a stronger preference for both types of anal sex, where the impact is especially strong for high-risk sex.

Estimates of parameters in the effort function (equation (10)) are reported in Table 15. All parameters are positive and significant. Positive parameters imply that the current-period relative payoff for anal sex (either moderate or high-risk sex) increases as more HIV− agents choose anal sex. This increase is relative to the flow utility of the low-risk option (no sex or oral sex only) along with the effort costs that HIV− men may face, both of which are normalized to zero. Recall that the effort function relates HIV+ men's choices to the proportion of HIV− men choosing anal sex weighted by the proportion of MSM who are HIV− at any given point in time. Therefore, if more men are HIV−, then the proportion choosing riskier sex will have a stronger impact on the behavior of HIV+ men. The model does not presuppose how HIV− men's behavior affects HIV+ men. For example, if HIV+

men were highly altruistic, they might refrain from riskier sex if more HIV− men entered the market. In that case, the parameters in the effort function would be negative. Instead, we find that as more HIV− men engage in moderate or high-risk sex, more HIV+ men engage in the same kinds of sex, even more so when HIV− men make up more of the market. This finding reflects increases in sex behavior among HIV+ men after HAART that cannot be explained by health improvements (see Table 6 and the accompanying discussion in Section 2.3).[41] Finally, we find a large, negative and significant scrap value of life. Given the normalization that low-risk sex utility is set to zero, the estimated parameter suggests that individual behavior is consistent with a preference for survival and low-risk sex over death.

To assess model fit, we compare model predictions of pre-HAART and post-HAART sex behavior decisions for each health state with analogous moments in our data. From the data, we calculate average sex behaviors for each health state before and after HAART. From the model, we also calculate average of sex behaviors, but weight the average by the observed proportion of observed state variables in each period. Results are in Figure 7, which shows that the model is able to capture differences in average sex behavior across health states along with changes to sex behavior after HAART is introduced. The model captures these changes despite our restriction that preferences remain stable over time. Instead the model generates changes in behavior through shifts in dynamic payoffs and through changes in behavior in response to market-level equilibrium effects.

The diffusion of HAART over time was not immediate and we capture this time variation in our model by explicitly conditioning on the proportion of HIV+ men using HAART in the health transition and seroconversion functions (equations (11) and (15), respectively). To emphasize this, we plot the simulated probability in each period that a white, college educated 35-year-old individual in each of the three health states chooses moderate-risk sex (see Figure 8). Looking at the figure, notice that HIV− men are more likely than HIV+ men to choose this option. In fact, they are more likely to choose anal sex in general. Further, for HIV− men, the dip in moderate-risk sex prior to HAART introduction is not due to the aging effect since the figure studies one age group. Rather, the decline reflects a riskier sex market as more MSM became HIV+ prior to HAART. After HAART, however, we see an increase from about 60% to about 67%. HIV+ men follow a similar pattern, reducing their sexual behavior prior to HAART. This is a response (via the effort function) to the reluctance of HIV− men to engage in risky sex prior to HAART. After HAART, we see HIV+ men increasing their choice of moderate-risk sex as more HIV− men became increasingly willing

[41] We note that the effort function is a mapping from HIV− men's optimal choices for sex to HIV+ men's flow utility. We cannot disentangle shifts due to altruism or fear of legal reprisal for infecting partners from effort costs of finding willing partners.

to engage in anal sex. HIV+ men with AIDS-level CD4 counts are less likely to have anal sex than the other groups since, as our model estimates show, agents in poor health obtain a lower relative flow utility from moderate and high-risk sex options. However, they also exhibit the largest upward shift after HAART is introduced. This follows from the estimated parameters of the effort function, which indicate that HIV+ men with low CD4 counts are more responsive to the optimal behavior of HIV− men in comparison to their healthier HIV+ counterparts.

6 Market Conditions and Equilibrium Sex Behavior

In this section we use the estimated model to investigate how optimal sex behavior changes as a function of market-level conditions. First, we consider HIV− and HIV+ agents separately and show how each set of agents best responds to market conditions and to the optimal behavior of other agents. Next, we show that market equilibrium is a fixed point where best response functions intersect.

6.1 HIV− Agents' Response to Market Conditions

In the model, infection rates are explicitly linked to the behavior and market saturation of HIV+ men. To illustrate this, we simulate the likelihood of seroconversion for moderate-risk sex as a function of the proportion of HIV+ men who are on HAART and who are engaged in either anal sex option. We fix the proportion of MSM who are HIV+ at the values in period 1 (1988), which is approximately 5%. Simulated infection rates for moderate-risk sex are plotted in Figure 9. The figure shows that infection rates increase as the number of HIV+ agents on HAART declines or when more HIV+ men choose anal sex.

How do HIV− men respond to this risk? We plot the proportion of uninfected men who optimally choose anal sex (either moderate or high-risk sex) under different market conditions in Figure 10. Several dynamics linking market conditions, preferences, forward-looking behavior and beliefs are summarized in this figure. Starting at the front corner of the figure: the proportion of HIV+ men choosing low-risk sex is set to zero and the proportion of HIV+ men on HAART is also set to zero. For a given market saturation of HIV+ men, this is the most dangerous sexual market for anal sex. In response, about 53% of uninfected men choose anal sex.

Suppose HIV+ men start using HAART. In response, HIV− men are more likely to engage in anal sex (move back from the front-right axis of Figure 10). If all HIV+ men are

on HAART, the risk of infection is nearly zero. The proportion of uninfected men choosing anal sex rises to about 73%. Still, a quarter of men, even in a market that is virtually free of the risk of HIV, do not choose anal sex. This proportion reflects the degree to which HIV− men respond to changing market quantities given their preferences. The magnitude of their response is identified from observed behavior changes in behavior after HAART in introduced. To fix ideas, suppose increases in risky behavior among HIV− men after HAART introduction had been larger. Then, the estimated model would have predicted a higher proportion of anal sex in a counterfactual market where infection is a zero probability event.

Returning to Figure 10, notice that as HIV+ men shift towards anal sex and if they are all on HAART (follow the plot in from the back-right to the back-left), there is virtually no impact on the behavior of HIV− men. Again, if all positive men are on HAART, they are nearly non-infectious. However, when we move along the axis from the back-right to the front-right, anal sex increases. As more HIV+ men choose low-risk sex, HIV− men shift towards anal sex. The logic is as follows. If no HIV+ men are on HAART, they are highly infectious. If all of them choose oral sex or no sex, then that option becomes riskier and therefore costlier. HIV− men, including those who prefer oral sex under lower infection rates, respond by substituting towards anal sex to avoid matching with HIV+ men. In other words, in a market where all HIV+ men choose the "safer" option, that option becomes relatively more risky and HIV− respond by switching to the now relatively more safe "riskier" option.

6.2 HIV+ Agents' Responses to Market Conditions

Next, we consider how HIV+ men optimally choose sex in light of the optimal behavior of uninfected men. Recall that the estimated parameters of the effort function indicate that the cost of effort for finding willing partners for HIV+ men who choose anal sex is higher when the optimal behavior of the HIV− men is to avoid anal sex. In conducting this simulation, we increase the proportion of HIV− men choosing anal sex, distributing them among moderate and high-risk sex according to relative choice probabilities implied by the model.

We plot the optimal responses of HIV+ men in Figures 11 and 12. First, we consider moderate-risk sex in Figure 11. As more HIV− men choose anal sex, more HIV+ men choose moderate-risk sex. This is a result of HIV+ men shifting away from the low-risk option as more partners are available for anal sex. The effect is slightly stronger when HIV− men constitute a larger portion of the total population of MSM. As rates of anal sex among HIV− men increase, however, the proportion of HIV+ men choosing moderate-risk sex declines. The reason can be seen in figure 12: HIV+ men substitute from the moderate-risk option to

the high-risk option.

6.3 Equilibrium in the Market for Sex among MSM

We now turn to our discussion of equilibrium behavior, defined as a fixed point where best-response functions intersect and illustrated in Figure 13. The dashed line in the figure plots the optimal response of HIV− men to their expectation of the sexual behavior of HIV+ men under the assumption that HAART has not been invented and where market weights are equal to those in 1988 (when about 5% of MSM were HIV+).[42] The line consisting of plus-signs plots the optimal behavior of HIV+ men to their expectations of HIV− men's behavior. For the optimal response of HIV+ men, however, we have switched the axes.

When the two lines do not intersect, HIV− or HIV+ men's expectations are inconsistent with the optimal responses of their counterparts on the market for sex. For example, suppose HIV− men expect HIV+ men to choose anal sex 40% of the time. Their best response is then to choose anal sex with approximately 70% probability (see Point A on Figure 13). However, if HIV− men are expected to choose anal sex with 70% probability, HIV+ men best respond by choosing anal sex with approximately 67% probability (and not 40%), so the market is not at equilibrium (see Point B on Figure 13). In equilibrium, behavior generates the same optimal response to which it is itself an optimal response. At the equilibrium point (Point C in Figure 13), HIV− men expect about 64% of HIV+ men to choose anal sex. They respond by choosing anal sex with approximately 68% probability. Conversely, if 68% of HIV− men choose anal sex, 64% of HIV+ choose risky sex, which is consistent with the expectations of HIV− men.

Figure 13 offers some lessons on the market for sex. First, the estimated slopes of the best response functions also determine market equilibrium. For example, if HAART use were more widespread, the elasticity of demand for sex in response to HIV+ men's sex behavior would be smaller in magnitude, which would raise equilibrium anal sex for both groups. There is also the issue of multiplicity of equilibria, which is ruled out by our estimates of the structural model. The figure shows that HIV− men's optimal choice of anal sex is a decreasing function of the optimal behavior of HIV+ men. In contrast, HIV+ men increase their anal sex behavior in response to HIV− men choosing anal sex. Given a decreasing best-response to an increasing best-response, if an equilibrium exists, it is unique, as can be seen in Figure 13. We note that the structure of the theoretical model does not guarantee uniqueness. For example, if HIV+ men best responded to increased sexual activity of HIV−

[42]This line is equivalent to the line in the surface plotted in Figure 10 extending from the bottom-right to the top-right axis.

men by *reducing* their sexual behavior, it would be possible that the best response functions would intersect more than once, implying the existence of multiple market equilibria. In other words, given the structure of the model, we can rule out multiplicity in this market using the estimated model parameters (identified using observed data).

7 The Value of Medical Innovation

In this section, we first use the estimated model to assess the value of HAART and then consider the value of a fully functional HIV vaccine. HAART lowered the probability of HIV infection and, conditional on HIV infection, reduced the probability of AIDS. By making both HIV and AIDS less likely, HAART not only increased expected lifespan, but also raised flow utility and lowered effort costs in each period. This can be seen in Tables 12, 14 and 15, which show health-state-dependent survival, utility and effort parameters, respectively. HAART-induced shifts in transition probabilities also brought changes in aggregate behavior, which further affected lifetime utility through market equilibrium effects. Since the estimated model explicitly captures the various ways that HAART introduction affected both HIV+ men and sexually active HIV− men at risk of infection, it is well-equipped for use in isolating the different ways that HAART created value. Throughout this section, our measure of value is willingness to accept payment (WTAP) in terms of expected remaining years of life.[43]

7.1 The Total Value of HAART

We begin by computing the total value of HAART, which includes its impact on infection rates and continuation payoffs (conditional on HIV infection) and accounts for equilibrium behavior changes among HIV+ and HIV− men who re-optimize in response to the innovation. First, we compute lifetime utility given HAART introduction, using estimated parameters and imposing equilibrium constraints. We use prevailing market conditions from 1995 where the number of HIV+ men without AIDS is 4.1% and with AIDS is 3.5%. Next, we assume that HAART use is 68% of HIV+ individuals, which is equivalent to average use about two years after it is introduced.

For each health state, we focus our attention on the value of HAART for an individual who is 35-years-old, white and college educated under these market conditions. We denote this vector of individual-level and market-level state variables as $\hat{Z}(1)$, where the 1 indicates HAART introduction. We compute choice-specific, net-of-disturbance value $\bar{V}_j(Z)$ given

[43]This measure is typically used in health economics when the model does not include income or a composite consumption good.

state variables $Z = \hat{Z}(1)$, where $\bar{V}_j(Z)$ is defined in equation (16). Next, we compute expected lifetime utility given state variables $\hat{Z}(1)$ assuming that agents maximize after drawing a choice-specific flow utility shock from an Extreme Value Type I distribution ϵ_j. Given distribution assumptions on ϵ_j, expected lifetime utility given state variables $\hat{Z}(1)$ is therefore given by:

$$\bar{V}\left(\hat{Z}(1)\right) = \gamma + \log\left[\sum_j \exp\left\{\bar{V}_j\left(\hat{Z}(1)\right)\right\}\right] \tag{21}$$

Similarly, we compute value assuming that HAART is not invented, imposing equilibrium constraints and permitting agents to re-optimize as before. Here, health transition probabilities and infection rates are set to pre-HAART values. We denote this value $\bar{V}(\hat{Z}(0))$, where the 0 denotes that HAART is not introduced.

Our next goal is to compute WTAP for HAART in expected years of life. To accomplish this, we relate computed lifetime utility to expected years of life. The procedure amounts to computing a compensating variation. Specifically, we increase value in the non-HAART counterfactual regime by increasing the likelihood of survival, assuming that individual behavior remains unchanged. The reason we impose that agents do not re-optimize given different survival rates is that we want to use an increase in expected lifespan as a commodity, i.e., as a device to measure value. We do not want to take into account how equilibrium changes in response to this increase in lifespan.[44] By raising the value of survival, we obtain a counterfactual lifetime value absent HAART and with survival increase Δs, which we denote $\bar{V}(\hat{Z}(0, \Delta s))$.

Next, we compute expected years of life for different values of Δs. We do this by using the estimated model for ages up to 75 and actuarial data on survival for men born in 1960 (who would be 35 in 1995) for higher ages. Notice that the model allows us to explicitly include the likelihood of survival for an MSM born in 1960 given the possibility of becoming infected with HIV. We repeat this exercise for 20 different counterfactual Δs and obtain 20 measures of lifetime utility. We then regress expected years of life on a second-order polynomial of computed lifetime utility. The estimated regression coefficients allow us to predict expected years of life for different levels of lifetime utility. Finally, we determine how many years of life correspond to HAART ($\bar{V}(\hat{Z}(1))$) versus the counterfactual scenario where HAART is not introduced ($\bar{V}(\hat{Z}(0))$).

Results on WTAP for HAART introduction measured in years of expected lifespan are found in the first two rows of Tables 16-18. In Table 16, we present valuation results for

[44] We explore more explicitly the role of equilibrium behavior changes in helping to determine the value of medical innovation below.

HIV− men. Absent HAART, we compute the expected lifespan for a 35-year-old college-aged white male in 1995 to be 65.26 years. This lifespan includes the counterfactual relatively high likelihood of becoming infected with HIV, which, absent HAART, would likely lead to AIDS and an early death. Next, we compute the additional years of life necessary to keep them indifferent (in terms of lifetime utility) to the scenario where HAART is invented. This occurs where expected lifespan is 72.44 years. For HIV− men, the value of HAART introduction thus corresponds to 7.19 years of life. In Tables 17 and 18, we compute the value of HAART for HIV+ men with high and low CD4 counts, respectively. For the first group, the corresponding value of HAART is 5.70 years of life. For the latter group, it is 4.22.

One immediate concern is that HIV− men appear to value HAART higher (in terms of expected remaining years of life) in comparison to HIV+ agents. Three factors explain this finding. First, HIV+ men have much lower expected lifespans. For HIV− men, 7.19 years means that HAART corresponds to about 24% of remaining lifespan of about 30 years. For HIV+ men who are healthy, WTAP is 5.70 years, corresponding to about 44% of remaining lifespan (about 13 years). For men with AIDS, 4.22 years means that HAART is worth about 45% of their remaining expected lifespan (about 10 years). Second, for HIV− men who expect to live longer than HIV+ men, additional expected years are more heavily discounted. Third, and as we will discuss in greater detail below, HAART lowered HIV infection rates for HIV− men, which benefits them directly; lower infection rates only affect HIV+ men through market equilibrium effects.[45]

7.2 The Value of HAART from Its Impact on Survival

An obvious benefit of HAART was its impact on the survival probability of HIV+ men. HAART had this effect by reducing the probability of getting AIDS and increasing the probability of recovering from AIDS. Isolating the value of HAART attributed to its impact on HIV+ survival occurs in two steps. First, we construct an innovation that improves HIV+

[45]Throughout this study, we assume that agents in the model randomly match on HIV-status. We make this assumption due to data limitations: we do not observe who sorts into high-risk or low-risk pools. It is possible that one subset of HIV− men in our sample sorts into a low-risk pool. For these men, we would over-estiamte the value of HAART. However, if one group of men sorts into a low-risk pool, another group must sort into a high-risk pool (relative to what we see in the data) in order to generate the infection rates we observe. We would under-estimate the value of HAART for men in the high-risk pool. The question is whether these two effects offset one another in a highly non-linear model of behavior. In results available from the authors, we test the sensitivity of our model to the random matching assumption. We show that if 10% of individuals in our sample sort into a group facing low-risk-sex infection rates across all choices (and the remaining 90% sort into a group facing enough of an elevated risk to generate the average infection rates we observe) the random matching assumption leads to a net 9% under-estimation of the value of HAART for HIV− men.

men's health as much as HAART did, but has no impact on HIV infection rates. We assume usage of this counterfactual drug to be 68% among HIV+ agents, the same rate of usage we assume when we value HAART. This counterfactual drug benefits individuals by reducing expected time spent in the AIDS-level, low-CD4 count health state ($h = 2$). Therefore, it improves survival, but also improves quality of life due to higher flow utility payoffs. The second step is to isolate the impact on survival by subtracting the expected utility gains for the incremental time spent in the non-AIDS state.[46]

Returning to Tables 16-18 (Row 3), we find that for HIV+ men, 77% of the total value of HAART is attributable to its effect on expected survival. HIV− men benefit from this effect as well since they anticipate some probability of becoming infected with HIV. For HIV− men, 18% of the total value of HAART is attributed to its impact on survival. Not surprisingly, the value of HAART from its impact on HIV+ survival is a major source of value for HIV+ men. It is less important for HIV− men, who benefit from this effect solely in expectation. Note that both calculations show that a considerable share of the value of HAART comes from sources other than its impact on HIV+ men's survival. We now examine these additional sources.

7.3 The Value of HAART from Additional Sources

One additional way that HAART created value was to improve quality of life conditional on HIV infection, since Table 14 shows that AIDS also implies lower utility parameters. In the previous section, we simulated a counterfactual drug that improved the health of HIV+ men, but that had no effect on infection rates. We now compute the value of this innovation. In doing so, we allow agents to re-optimize their sexual behavior and also impose that equilibrium constraints hold. Computing the value of this drug is similar to our previous calculation to isolate survival. The difference is that we skip the second step and therefore do not limit our attention to survival. Returning to Tables 16-18 (Row 4) we find that when we take account of both survival and higher utility given HIV infection, the proportion of the value of HAART we can account for rises to 84-85% for HIV+ men and to 20% for HIV− men. This means that for HIV− individuals, 20% of the total value of HAART is its option value. It embodies the payoff to HIV− men for a counterfactual medical innovation that shifts HIV+ health state transition probabilities, thereby improving survival and quality of life conditional on HIV infection. Notice, in computing the option value of HAART, we have not performed a decomposition. Rather, we have studied the value of a counterfactual treatment with only one of the effects of HAART. Calculating value in this way allows us to

[46] As this step is fairly complex, we relegate details to Appendix D.

bypass how decompositions are not unique to the order of the decomposition. It also permits a more realistic view of how a drug with only one of the effects of HAART would create value as we recalculate optimal behavior in light of general equilibrium effects.[47]

7.4 Equilibrium Effects and the Value of Medical Innovation

In this section we highlight the the role of market equilibrium effects in determining the value of HAART. To accomplish this, we impose that market-level behavior remains fixed at pre-HAART levels, but allow individual agents to re-optimize and then re-compute the expected lifetime utility. In other words, we allow individual agents to re-optimize, but do not impose that equilibrium constraints hold. To fix ideas, if we hold HIV− men's behavior fixed at pre-HAART quantities, this is only from the perspective of HIV+ men. In Table 19, we show value under this counterfactual. First, holding HIV− behavior fixed lowers the value of HAART by 2% for HIV− men. This might seem to contradict the idea that agents are allowed to re-optimize and that holding behavior fixed should only affect HIV+ men. However, recall that HIV− men anticipate some possibility of becoming HIV+. Holding HIV− behavior fixed, which is costly for HIV+ men, is also costly, in expectation, for HIV− men, reducing the value of HAART by 2%. When we consider equilibrium effects on HIV+ men's valuation of HAART, we see that about 20% of the total value of HAART is lost if HIV− men do not increase their sexual behavior. An important benefit of HAART for HIV+ men is therefore its impact HIV+ men's quality of life resulting from increases in sexual behavior among HIV− men.

Next, when HIV+ equilibrium behavior is fixed at pre-HAART levels, the value of HAART rises by about 4% for HIV+ men. Prior to HAART, HIV+ men had less sex, making the market safer for HIV− men. About 4% of the value of HAART is therefore lost since HIV+ men increase their risky behavior. This loss is fairly small and reflects how HIV is a less costly and less likely condition after HAART is introduced.[48] However, this is not to say that infected men's behavior has no impact on the HIV− men's welfare. To make this point, we compute lifetime utility of HIV− men under a counterfactual regime without HAART, but where HIV+ men engage in their post-HAART optimal behavior. HIV−

[47]Relatedly, valuing each of the effects of HAART by simulating counterfactual treatments, each of which has only one of the effects of HAART, means that we can account for the fact that the effects of HAART can be complements or substitutes. As such, the values of each effect do not necessarily sum to the total value of HAART. For HIV− men, for example, a drug that lowers infection rates and a drug that improves health conditional on infection are substitutable goods. Therefore, as we show in supplementary calculations, from the perspective of HIV− men, the sum of the value of counterfactual innovations mimicking each effect of HAART is larger than the total value of HAART.

[48]By construction, it has no impact on the value of HAART for HIV+ men.

men's WTAP for such a regime is -1.58 years, corresponding to a loss of lifetime utility that is equivalent to about 22% of the total value of HAART. This quantity reflects how HIV infection is far more likely absent the attenuating effect of HAART on infection and is still very costly absent effective treatment.

7.5 The Value of Fully Functional Vaccine

Having discussed the various ways that HAART created value, we now use our estimated model to assess the value of a counterfactual, fully functional vaccine. To compute this quantity, we introduce in place of HAART a counterfactual medical innovation that reduces infection to nearly 0%.[49] In Tables 16-18 (Row 5), we see that, for HIV− men, the value of a counterfactual vaccine exceeds the value of HAART: 9.07 years versus 7.19 years. 9.07 years is quite large and underscores the seriousness of the AIDS epidemic. On the other hand, the value of a vaccine is only about 26% larger than the value of HAART. The limited incremental value of a vaccine reflects two key factors. First, HAART already functions as a highly effective (if imperfect) vaccine that lowers infection. Second, since HAART is a medication that vastly lowers the cost of HIV infection, a vaccine preventing HIV is of less value.[50]

A vaccine is also valuable for HIV+ men, who would value it roughly 9% higher than HAART. Faced with a vaccine that eliminates risk of HIV infection, HIV− men would increase their riskier sexual behavior. This shift in behavior would lower the effort costs of finding partners willing to engage in riskier sexual behaviors, which would benefit HIV+ men. Therefore, HIV+ men would be willing to accept payment for a vaccine they would never use. This finding is similar to the idea of the option value of HAART from the perspective of HIV− men. Since it captures dynamics, our model can be used to assess the value of a drug used by HIV+ men, but that also benefits HIV− men in expectation; since it captures market equilibrium effects, our model can also be used to assess the value of a vaccine used solely by HIV− men from the perspective of HIV+ men, who benefit due to resulting shifts in equilibrium behavior. In the case of infectious diseases like HIV, we have shown that

[49]Here, we do not set infection rates to zero since our mapping from market level quantities to seroconversion rates indicate that infection rates when all agents are on HAART or where no HIV+ agents are engaged in risky sex puts seroconversion rates at tiny, though positive numbers. These rates likely capture, for example, transmission through other means (e.g., intravenous drug use) or measurement error. We set infection probability to zero in supplementary analysis available from the authors, though results are largely unchanged. Our estimates likely represent an upper bound on the value of a vaccine since it is unlikely that a vaccine would fully functional. Moreover, estimates of the value of the vaccine for HIV+ men assume universal uptake among HIV−, which is also unlikely.

[50]In this sense, a cure for HIV and a fully functional vaccine for HIV function are perfect substitutes from the perspective of HIV− men.

assessing new medical technologies, including how they affect behavior and create value, requires consideration of dynamic decision-making and market equilibrium effects.

8 Conclusion

In this paper, we have developed a framework for measuring the value of medical innovation in the context of infectious disease. We have applied our framework to HIV to assess the value of HAART and a counterfactual, fully functional vaccine. Our framework captures several ways that medical innovation creates value in the context of infectious disease. In concluding, we highlight two of these effects: dynamics and general equilibrium effects. First, our framework accounts for the option value of new medical technologies. The idea is that agents who are not infected, but who anticipate some probability of future infection, benefit in expectation from the provision of a new medical technology they are unlikely to ever consume. This is an important consideration in computing value since, in most medical contexts, the number of individuals who may eventually become infected far exceeds the number who actually do fall ill. As we show, the option value of HAART is less than the total value of HAART for HIV− and HIV+ men. However, even among groups exhibiting high rates of HIV infection, e.g., MSM, the option value accrues to a much larger group.

The framework we have developed also captures the role of market-level shifts in equilibrium behavior in determining how new medical technologies affect behavior and create value. In the context of infectious disease, market equilibrium effects imply additional costs and benefits arising from new medical technologies. For example, a benefit of HAART from the perspective of infected agents is that it led to disinhibition on the part of uninfected men, which lowered the effort costs of finding partners willing to engage in risky sex. From the perspective of uninfected men, HAART increased risky behavior among HIV+ men, which *ceteris paribus* raised infection rates. Importantly, our model captures these market equilibrium effects within a single dynamic framework so that we can draw precise conclusions about their impact relative to the total value of the innovation. For example, we show that increased infection rates after HAART due to the increased sex behavior of HIV+ agents certainly attenuated the total value of HAART for HIV− men, but only slightly. Therefore, we can use our estimated model to conclude with confidence that HAART generated great value for HIV+ and HIV− men alike.

References

Adda, Jérôme. 2007. "Behavior Towards Health Risks: An empirical Study Using the "Mad Cow" Crisis as an Experiment." *Journal of Risk and Uncertainty* 35 (3):285–305.

Ahituv, Avner, V Joseph Hotz, and Tomas Philipson. 1996. "The Responsiveness of the Demand for Condoms to the Local Prevalence of AIDS." *Journal of Human Resources* 31 (4):869–897.

Arcidiacono, P., A. Khwaja, and L. Ouyang. 2012. "Habit Persistence and Teen Sex: Could Increased Access to Contraception Have Unintended Consequences for Teen Pregnancies?" *Journal of Business and Economic Statistics* 30 (2):312–325.

Arcidiacono, P., H. Sieg, and F. Sloan. 2007. "Living Rationally under the Volcano? Heavy Drinking and Smoking among the Elderly." *International Economic Review* 48 (1):37–65.

Arcidiacono, Peter, Andrew W Beauchamp, and Marjorie B McElroy. 2010. "Terms of Endearment: An Equilibrium Model of Sex and Matching." NBER working paper.

Attia, Suzanna, Matthias Egger, Monika Müller, Marcel Zwahlen, and Nicola Low. 2009. "Sexual Transmission of HIV According to Viral Load and Antiretroviral Therapy: Systematic Review and Meta-Analysis." *AIDS* 23 (11):1397–1404.

Auld, M.C. 2003. "Choices, Beliefs, and Infectious Disease Dynamics." *Journal of Health Economics* 22 (3):361–377.

———. 2006. "Estimating Behavioral Response to the AIDS Epidemic." *Contributions to Economic Analysis & Policy* 5 (1):1235–1235.

Austen-Smith, David and Roland G. Fryer. 2005. "An Economic Analysis of 'Acting White'." *Quarterly Journal of Economics* 120 (2):551–583.

Bechange, Stevens, Rebecca Bunnell, Anna Awor, David Moore, Rachel King, Jonathan Mermin, Jordan Tappero, Kenneth Khana, and Bradford Bartholow. 2010. "Two-Year Follow-Up of Sexual Behavior among HIV-Uninfected Household Members of Adults Taking Antiretroviral Therapy in Uganda: No Evidence of Disinhibition." *AIDS and Behavior* 14 (4):816–823.

Becker, Gary S. and Kevin M. Murphy. 1988. "A Theory of Rational Addiction." *Journal of Political Economy* 96 (4):675–700.

Bhaskaran, K, O Hamouda, M Sannes, and et al. 2008. "Changes in the Risk of Death after HIV Seroconversion Compared with Mortality in the General Population." *Journal of the American Medical Association* 300 (1):51–59.

Cawley, J. and C.J. Ruhm. 2012. "The Economics of Risky Health Behaviors." In *Handbook of Health Economics*, vol. 2, edited by T. G. McGuire, M.V. Pauly, and P.P. Barros, chap. 3. Elsevier: New York, 95–199.

Dilley, James W, William J Woods, and William McFarland. 1997. "Are Advances in Treatment Changing Views about High-Risk Sex?" *New England Journal of Medicine* 337 (7):501–502.

Fryer, Roland. 2003. "An Economic Approach to Cultural Capital." Mimeo, University of Chicago.

Gersovitz, Mark, Hanan G Jacoby, F Seri Dedy, and A Gozé Tapé. 1998. "The Balance of Self-Reported Heterosexual Activity in KAP Surveys and the AIDS Epidemic in Africa." *Journal of the American Statistical Association* 93 (443):875–883.

Kaslow, RA and DG Ostrow. 1987. "The Multicenter AIDS Cohort Study (MACS): Rationale, Organization and Selected Characteristics of the Participants." *American Journal of Epidemiology* 126:310–318.

Keane, M.P. and K.I. Wolpin. 2002. "Estimating Welfare Effects Consistent with Forward-Looking Behavior. Part I: Lessons from a Simulation Exercise." *Journal of Human Resources* 37 (3):570–599.

Koszegi, Botond and Jonathan Gruber. 2001. "Is Addiction "Rational"? Theory and Evidence." *Quarterly Journal of Economics* 116 (4):1261–1303.

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

Lakdawalla, D., N. Sood, and D. Goldman. 2006. "HIV Breakthroughs and Risky Sexual Behavior." *Quarterly Journal of Economics* 121 (3):1063–1102.

Lamourelle, G., P. Harrison, J. Rowley, and M Warren. 2005. "Tracking Funding for Preventive HIV Vaccine Research and Development." Tech. rep., HIV Vaccines and Microbicides Resource Tracking Working Group.

McGinnis, J.M. and W.H. Foege. 1993. "Actual Causes of Death in the United States." *Journal of the American Medical Association* 270 (18):2207–2212.

Mechoulan, Stephane. 2007. "Risky Sexual Behavior, Testing, and HIV Treatments." *Forum for Health Economics and Policy* 10 (2):1–51.

Ostrow, DG, A. Monjan, J. Joseph, M. VanRaden, R. Fox, L. Kingsley, J. Dudley, and J. Phair. 1989. "HIV-Related Symptoms and Psychological Functioning in a Cohort of Homosexual Men." *American Journal of Psychiatry* 146 (6):737.

Papageorge, N.W. 2015. "How Does Medical Innovation Create Value? Health, Human Capital and the Labor Market." Mimeo, Johns Hopkins University.

Paula, Áureo De, Gil Shapira, and Petra E Todd. 2014. "How Beliefs about HIV Status Affect Risky Behaviors: Evidence from Malawi." *Journal of Applied Econometrics* 29 (6):944–964.

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

Stuber, Jennifer, Sandro Galea, and Bruce G Link. 2008. "Smoking and the Emergence of a Stigmatized Social Status." *Social Science & Medicine* 67 (3):420–430.

Thaler, Richard and Sherwin Rosen. 1976. "The Value of Saving a Life: Evidence from the Labor Market." In *Household Production and Consumption*. NBER, 265–302.

Thornton, R.L. 2008. "The Demand for, and Impact of, Learning HIV Status." *American Economic Review* 98 (5):1829–1863.

Van der Klaauw, W. and K.I. Wolpin. 2008. "Social Security and the Retirement and Savings Behavior of Low-Income Households." *Journal of Econometrics* 145 (1):21–42.

9 Tables and Figures

Table 1: Summary Statistics — Individual Choices and States

	Full Sample	HIV+ (Baseline)	HIV−	Sero-converters
	N=2,426	N=1,268	N=766	N=392
HIV Status	1.00	0.52	0.32	0.16
Age (mid-1989)	37.01	36.32	38.69	35.97
$\geq$ College Degree	0.57	0.53	0.65	0.53
Non-White	0.18	0.22	0.13	0.15
Anonymous Sex before mid-1989	0.35	0.35	0.34	0.35
Average CD4 (Pre-HAART)	640.84	404.11	1045.50	569.20
Average CD4 (Post-HAART)	696.13	479.48	976.69	536.48
Alive at Sample End	0.58	0.43	0.90	0.58
Given Survival until HAART	0.89	0.83	0.97	0.83

The subsample for analysis is a panel of 2,426 individuals over 27 semi-annual periods (mid-1989 to mid-2002). 2,107 joined the sample in 1984. 319 joined in a second cohort arriving between 1987-1989. Both groups therefore arrived prior to the start of our analysis sample period. Individuals are observed an average of 12 times (ranging from 1-27 times). The total number of observations is 30,274. All entries except for age and CD4 count are proportions within infection categories. Categories are: HIV+ at baseline data collection (1984), HIV− throughout data collection period and observed seroconverter, i.e., transition from negative to positive during sample collection period.

Table 2: Summary Statistics — Market-Level Quantities

	P[HIV+ | MSM]			P[On HAART | HIV+ & MSM]		
Year	All	High CD4	Low CD4	All	High CD4	Low CD4
1989	0.042	0.028	0.014	0.000	0.000	0.000
1990	0.046	0.031	0.016	0.000	0.000	0.000
1991	0.053	0.034	0.018	0.000	0.000	0.000
1992	0.068	0.042	0.026	0.000	0.000	0.000
1993	0.070	0.043	0.027	0.000	0.000	0.000
1994	0.072	0.041	0.032	0.000	0.000	0.000
1995	0.076	0.041	0.035	0.007	0.000	0.015
1996	0.078	0.047	0.031	0.287	0.194	0.427
1997	0.079	0.052	0.027	0.574	0.518	0.683
1998	0.082	0.063	0.019	0.657	0.659	0.650
1999	0.085	0.065	0.020	0.676	0.691	0.624
2000	0.088	0.066	0.021	0.690	0.716	0.607
2001	0.092	0.070	0.022	0.692	0.702	0.657
2002	0.099	0.082	0.017	0.687	0.684	0.701

This table summarizes market-level quantities used in our analysis. These quantities capture agent beliefs about disease dynamics for the MSM partner market. We use data from city and county health departments along with CDC data and MACS data to calculate the probability that an MSM sex partner is HIV+ with a high or low CD4 count. Details on the calculations are discussed in Table A1 in Appendix A. To calculate the proportion of HIV+ individuals using HAART, we use MACS data.

Table 3: Summary Statistics — Transitions among Health States

	Time $t+1$ Health State			
	HIV−	HIV+ High CD4	HIV+ Low CD4	Death
Time t Health State				
Pre-HAART				
HIV−	0.981	0.012	0.000	0.007
HIV+ & High CD4	0.000	0.870	0.121	0.009
HIV+ & Low CD4	0.000	0.100	0.769	0.131
Post-HAART				
HIV−	0.994	0.003	0.000	0.004
HIV+ & High CD4	0.000	0.923	0.073	0.004
HIV+ & Low CD4	0.000	0.244	0.716	0.040

Unconditional health state to state transitions for the entire sample period and for the pre- and post-HAART eras.

Table 4: Summary Statistics — Sex Behavior Decisions

	Low Risk	Mod. Risk	High Risk
All Subjects	0.42	0.53	0.05
HIV−	0.38	0.57	0.05
HIV+ & High CD4	0.39	0.56	0.05
HIV+ & Low CD4	0.57	0.39	0.03
Pre-HAART	0.42	0.54	0.04
HIV−	0.38	0.58	0.04
HIV+ & High CD4	0.38	0.58	0.04
HIV+ & Low CD4	0.59	0.38	0.03
Post-HAART	0.41	0.52	0.06
HIV−	0.39	0.55	0.06
HIV+ & High CD4	0.39	0.53	0.07
HIV+ & Low CD4	0.54	0.41	0.04
Not Anonymous	0.42	0.54	0.04
Anonymous	0.42	0.52	0.06

In each period, agents choose one of three sex options (i) low-risk sex (no intercourse or oral intercourse only), (ii) moderate-risk sex (anal sex with a single partner or with multiple partners where condoms are used with at least one of the partners consistently) or (iii) high-risk sex (anal sex with multiple partners where condoms are not used consistently with any partner.)

Table 5: Summary Statistics — Persistence in Sex Choices

	Time $t+1$ Sex Choice		
	Low Risk	Mod. Risk	High Risk
Time t Sex Choice			
All Observations			
Low-Risk Sex	0.80	0.19	0.01
Moderate-Risk Sex	0.16	0.79	0.05
High-Risk Sex	0.09	0.57	0.34
Pre-HAART			
Low-Risk Sex	0.79	0.20	0.01
Moderate-Risk Sex	0.17	0.79	0.05
High-Risk Sex	0.10	0.62	0.28
Post-HAART			
Low-Risk Sex	0.81	0.17	0.01
Moderate-Risk Sex	0.14	0.80	0.06
High-Risk Sex	0.08	0.52	0.40

Period $t+1$ sex choices conditional on period t tex choices for all observations, for the pre-HAART era and for the post-HAART era.

Table 6: Preliminary Analysis — Sexual Behavior

	Moderate-Risk Sex		High-Risk Sex	
Variable	Coeff.	Std. errors	Coeff.	Std. errors
HIV−	2.61	0.29	-0.76	0.51
Post-HAART × HIV−	0.29	0.08	0.54	0.16
HIV+ & high CD4	2.55	0.31	-0.69	0.46
Post-HAART × HIV+ & High CD4	0.32	0.08	0.85	0.15
HIV+ & low CD4	1.01	0.36	-1.84	0.75
Post-HAART × HIV+ & Low CD4	0.34	0.11	0.60	0.24
Age × HIV−	-0.06	0.01	-0.04	0.01
Age × HIV+ & High CD4	-0.06	0.01	-0.04	0.01
Age × HIV+ & Low CD4	-0.06	0.01	-0.05	0.02
College Educated	0.14	0.08	-0.04	0.13
Engaged in Anon. Sex	0.08	0.08	0.46	0.13
Non-White	0.19	0.10	-0.09	0.16
Continous CD4/100 × HIV+ & High CD4	0.03	0.02	<0.00	0.03
Continous CD4/100 × HIV+ & Low CD4	0.47	0.05	0.42	0.11

Multinomial logistic regression explaining sexual behavior as a function of health interacted with an indicator for the post-HAART era along with other sociodemographic measures.

Table 7: Preliminary Analysis — Sexual Behavior for HIV− Individuals

	Moderate-Risk Sex		High-Risk Sex	
Variable	Coeff.	Std. errors	Coeff.	Std. errors
Prob. that HIV+ Use HAART	0.46	0.13	0.92	0.26
Age	-0.06	0.01	-0.04	0.01
College Educated	0.01	0.12	-0.19	0.21
Engaged in Anon. Sex	0.19	0.12	0.52	0.21
Non-White	0.16	0.17	0.02	0.26
Intercept	2.68	0.29	-0.67	0.53

Multinomial logistic regression explaining risky sexual behavior for HIV− agents as a function of technology diffusion (measured by proportion of HIV+ agents using HAART at time t) and sociodemographic variables.

Table 8: Preliminary Analysis — Serconversion

	[1]	[2]	[3]	[4]
Moderate Risk Sex	1.25	1.25	1.31	1.32
	(0.31)	(0.31)	(0.31)	(0.31)
High Risk Sex	1.78	1.79	1.88	1.87
	(0.43)	(0.43)	(0.43)	(0.43)
Age	-0.06	-0.06	-0.03	-0.03
	(0.01)	(0.02)	(0.02)	(0.02)
College Educated	-0.37	-0.37	-0.42	-0.42
	(0.22)	(0.22)	(0.21)	(0.21)
Engaged in Anon. Sex	0.32	0.32	0.25	0.24
	(0.22)	(0.22)	(0.22)	(0.22)
Non-White	0.26	0.26	0.33	0.34
	(0.25)	(0.25)	(0.25)	(0.25)
Avg. Anal - HIV+	.	-2.77	.	4.89
		(3.64)		(3.96)
Prob. that HIV+ Use HAART	.	.	-1.87	-2.07
			(0.5)	(0.54)
Observations	12908	12908	12908	12908

Logistic regression explaining seroconversion as a function of sexual behavior, sociodemographic variables and the decisions of HIV+ agents, including HAART use and sex behavior.

Table 9: Structural Parameter Estimates — HAART Choice

	High CD4			Low CD4		
Variables	Coefficients	Std. Errors	Marg. Effects	Coefficients	Std. errors	Marg. Effects
1994.5	.	.	.	-5.73	1.04	0.01
1995	.	.	.	-4.43	0.59	0.02
1995.5	-2.92	0.38	0.04	-2.79	0.39	0.09
1996	-1.03	0.25	0.22	-0.51	0.35	0.46
1996.5	-0.15	0.24	0.39	0.33	0.35	0.64
1997	0.52	0.24	0.52	0.54	0.36	0.68
1997.5	1.00	0.24	0.62	0.67	0.37	0.70
1998	1.11	0.24	0.64	0.38	0.38	0.64
1998.5	1.10	0.24	0.64	0.34	0.37	0.64
1999	1.28	0.25	0.67	0.26	0.39	0.62
1999.5	1.12	0.25	0.64	0.72	0.39	0.71
2000	1.41	0.26	0.69	0.23	0.39	0.62
2000.5	1.34	0.26	0.68	0.53	0.39	0.68
2001	1.34	0.26	0.68	0.42	0.40	0.66
2001.5	1.27	0.26	0.67	0.45	0.42	0.66
2002	1.26	0.27	0.67	0.61	0.44	0.69
2002.5	1.24	0.26	0.66	0.56	0.43	0.68
Age	-0.01	0.00	-0.001	0.01	0.01	0.001
College	0.23	0.06	0.04	0.21	0.11	0.04
Anon.	-0.13	0.07	-0.02	-0.28	0.10	-0.05
Non-White	-0.56	0.08	-0.11	-0.50	0.14	-0.10

Structural model estimates: HAART choice coefficients from equation (25). For each variable, marginal effects are calculated at the means of other variables.

Table 10: Structural Parameters — High to Low CD4

Variables	Coefficients	Std. errors
Age	0.01	0.00
College	-0.05	0.06
Anon.	0.02	0.06
Non-White	-0.02	0.08
Using HAART	-0.58	0.09
Constant	-2.57	0.17

Structural model estimates: Health transitions for HIV+ (high to low CD4) from equation (24).

Table 11: Structural Parameter Estimates — Low to High CD4

Variables	Coefficients	Std. errors
Age	0.01	0.00
College	0.02	0.04
Anon.	0.05	0.04
Non-White	0.26	0.05
Using HAART	0.72	0.06
Constant	-2.47	0.10

Structural model estimates: Health transitions for HIV+ (low to high CD4) from equation (24).

Table 12: Structural Parameters — Survival

Variables	Coefficients	Std. errors
HIV+ & High CD4	0.12	0.17
HIV+ & Low CD4	2.98	0.13
Age	0.00	0.00
College	-0.24	0.08
Non-White	-0.04	0.11
Constant	-5.20	0.26

Structural model estimates: Survival function parameters from equation (23).

Table 13: Structural Parameters — Seroconversion

Variables	Coefficients	Std. errors
Low-Risk Sex (Constant)	-7.87	1.09
Low-Risk Sex (Slope)	35.70	17.83
Moderate-Risk Sex (Constant)	-5.62	0.43
Moderate-Risk Sex (Slope)	40.79	14.12
High-Risk Sex (Constant)	-4.92	1.03
High-Risk Sex (Slope)	34.37	36.08

Structural model parameters: Mappings from market quanities to seroconversion, corresponding to the Φ in equation (15).

Table 14: Structural Parameters — Utility

Variables	Safer Anal		Riskier Anal	
	Coefficients	Std. errors	Coefficients	Std. errors
HIV−	3.14	0.14	-0.10	0.29
HIV+ & High CD4	2.76	0.12	-1.31	0.30
HIV+ & Low CD4	1.74	0.18	-1.41	0.50
HIV− × Age	-0.05	0.00	-0.02	0.01
HIV+ & High CD4 × Age	-0.06	0.00	-0.02	0.01
HIV+ & Low CD4 × Age	-0.05	0.00	-0.03	0.01
College	0.14	0.03	-0.06	0.06
Anonymous Sex	0.06	0.03	0.37	0.06
Non-White	0.20	0.04	-0.02	0.08

Structural model estimates: Utility parameters from equation (9).

Table 15: Structural Parameters — Effort and Scrap Value

Variables	Coefficients	Std. errors
HIV+ & High CD4 & Moderate-Risk Sex	0.01	0.00
HIV+ & High CD4 & High-Risk Sex	0.03	0.01
HIV+ & Low CD4 & Moderate-Risk Sex	0.03	0.01
HIV+ & Low CD4 & High-Risk Sex	0.03	0.01
Scrap Value of Life	-140.21	24.77

Structural model estimates: Effort function parameters from equation (10) along with the scrap value.

Table 16: Value of HAART (HIV−)

Factual and Counterfactual Innovations	Valued in Units of Expected Longevity		
	Lifespan	Δ Years	% of Δ
HAART is invented (as observed)	72.44	7.19	1.00
HAART is not invented	65.26	0.00	0.00
HAART affects HIV+ survival only	66.55	1.29	0.18
HAART affects HIV+ transition probabilities	66.69	1.43	0.20
A vaccine ends seroconversion	74.32	9.07	1.26

Value calculations for different components of HAART for HIV− men.

Table 17: Value of HAART (HIV+ and High CD4)

Factual and Counterfactual Innovations	Valued in Units of Expected Longevity		
	Lifespan	Δ Years	% of Δ
HAART is invented (as observed)	53.62	5.70	1.00
HAART is not invented	47.92	0.00	0.00
HAART affects HIV+ survival only	52.28	4.37	0.77
HAART affects HIV+ transition probabilities	52.71	4.79	0.84
A vaccine ends seroconversion	54.13	6.21	1.09

Value calculations for different components of HAART for HIV+ men with high CD4 counts.

Table 18: Value of HAART (HIV+ and Low CD4)

Factual and Counterfactual Innovations	Valued in Units of Expected Longevity		
	Lifespan	Δ Years	% of Δ
HAART is invented (as observed)	48.42	4.22	1.00
HAART is not invented	44.21	0.00	0.00
HAART affects HIV+ survival only	47.45	3.24	0.77
HAART affects HIV+ transition probabilities	47.78	3.57	0.85
A vaccine ends seroconversion	48.79	4.58	1.09

Value calculations for different components of HAART for HIV+ men with low CD4 counts.

Table 19: Value of HAART: Equilibrium Effects

Value For	Counterfactual Market Equilibrium Conditions	Valued in % of Total Value of HAART
HIV−	HIV− Choices: Pre-HAART Levels	0.98
HIV−	HIV+ Choices: Pre-HAART Levels	1.04
HIV−	All Choices: Pre-HAART Levels	1.02
HIV−	HIV+ Choices: Post-HAART Levels (in Pre-Era)	-0.22
HIV+ High CD4	HIV− Choices: Pre-HAART Levels	0.80
HIV+ Low CD4	HIV− Choices: Pre-HAART Levels	0.81

Value calculations for different components of HAART, isolating the role of equilibrium effects.

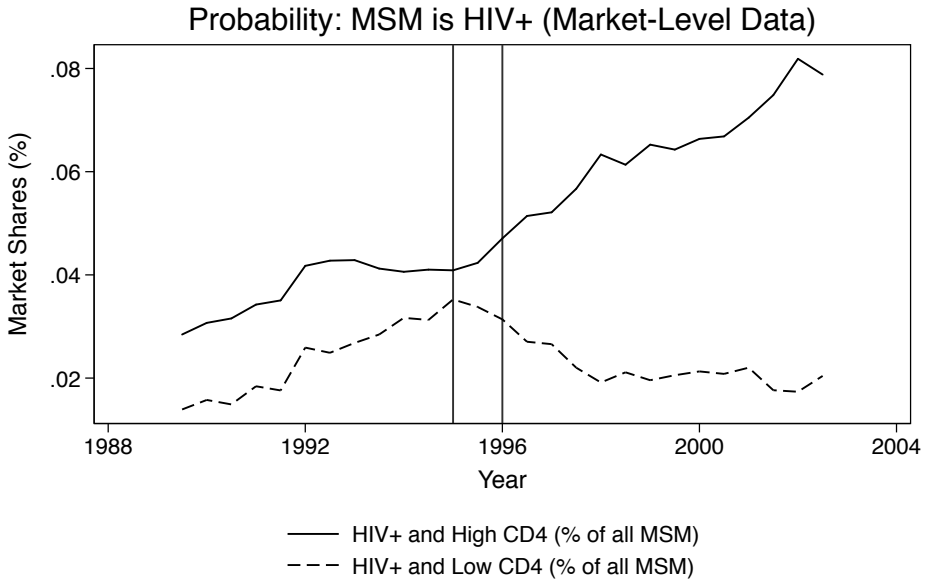

Figure 1: Market level data: the proportion of MSM who are HIV+ with high and low CD4 counts (constructed using national and county level data along with MACS data).

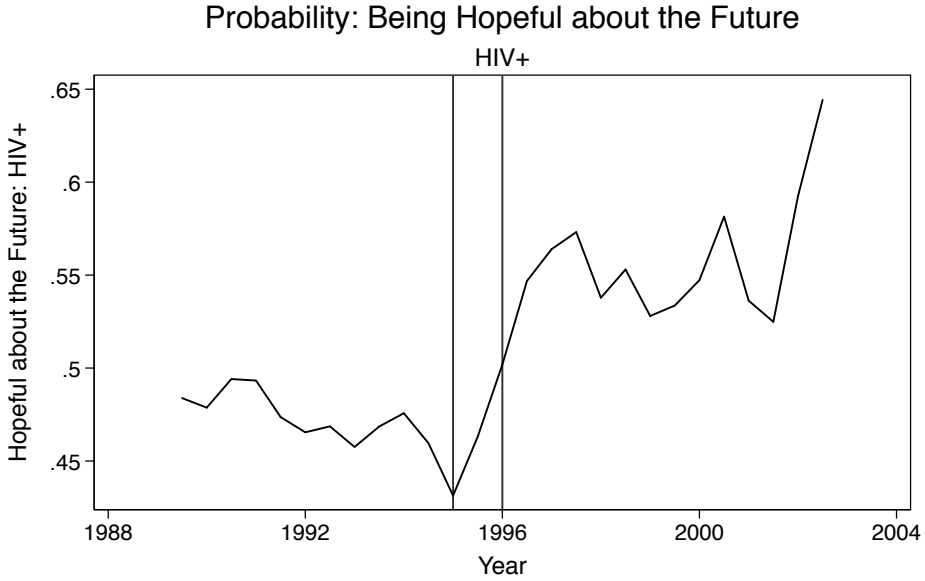

Figure 2: Portion of HIV+ agents who are hopeful about the future most or all of the time in the week prior to their visit.

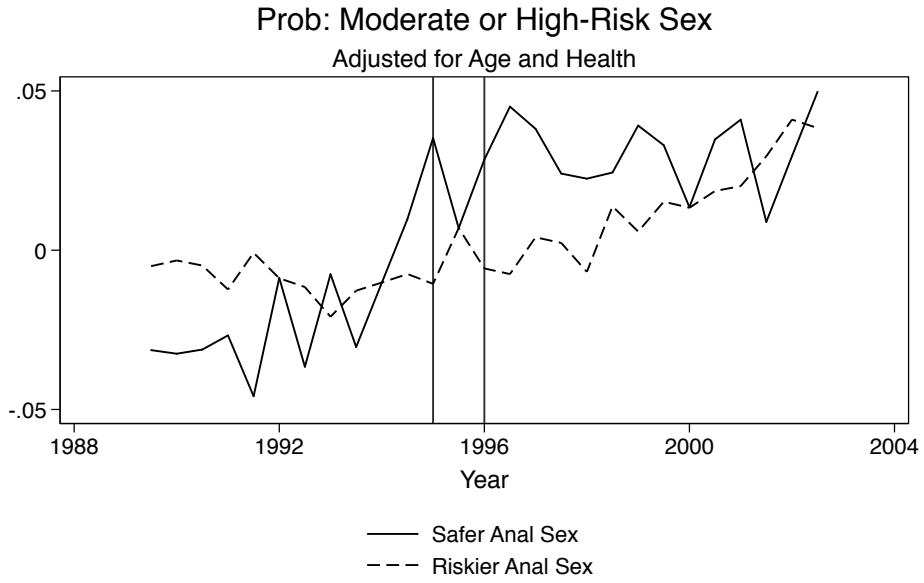

Figure 3: Probability of moderate or high-risk sex over time, corrected for health and age.

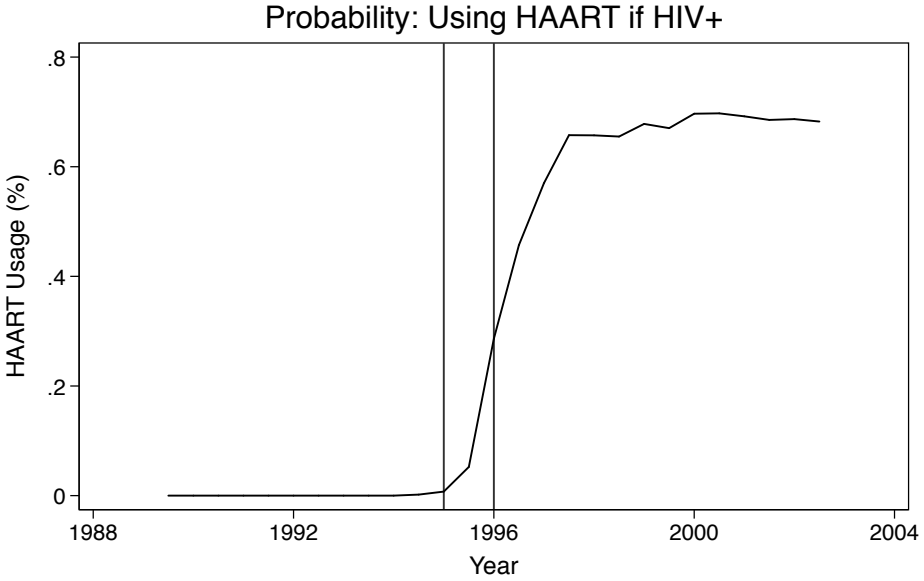

Figure 4: Diffusion of a new medical technology: uptake of HAART among HIV+ men over time.

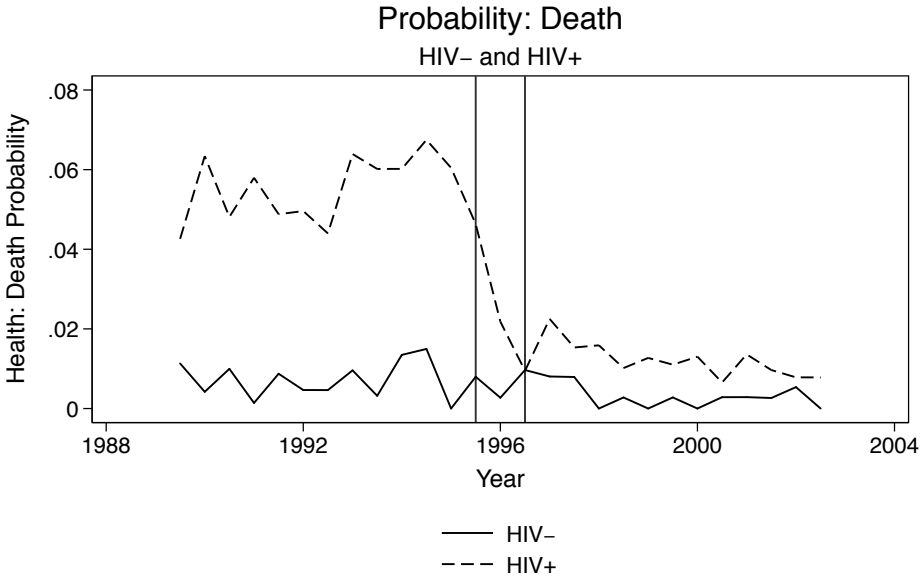

Figure 5: Death probability over time.

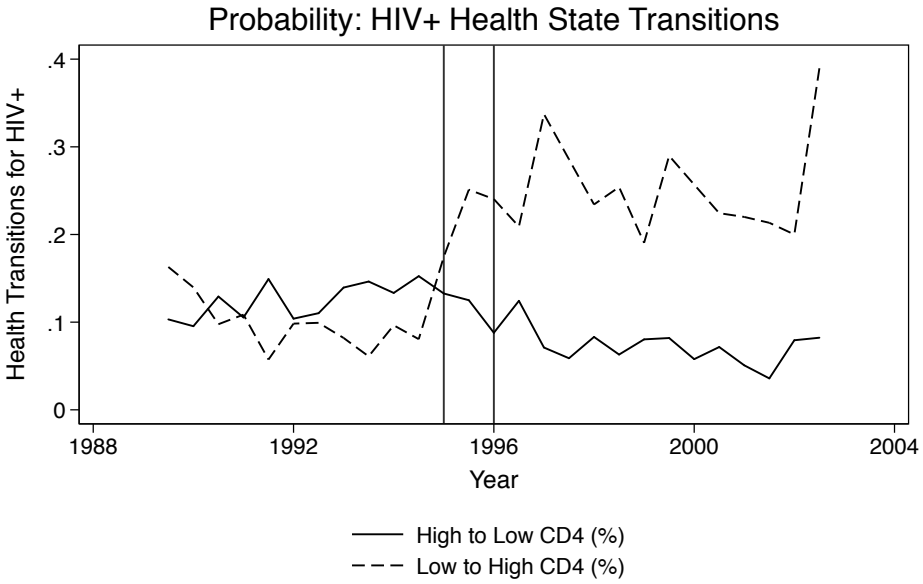

Figure 6: AIDS and recovery from AIDS over time for HIV+ men.

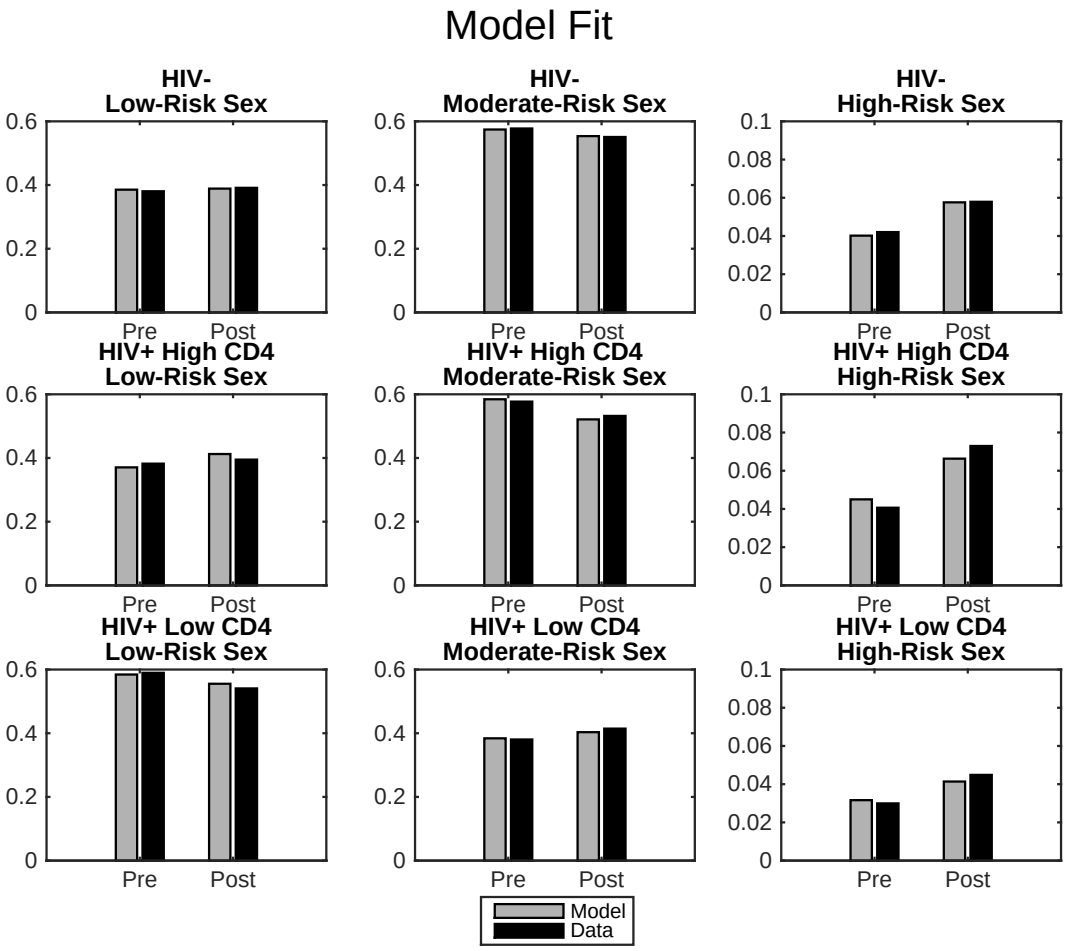

Figure 7: Average sex behavior by health state in the sample and predicted by the model.

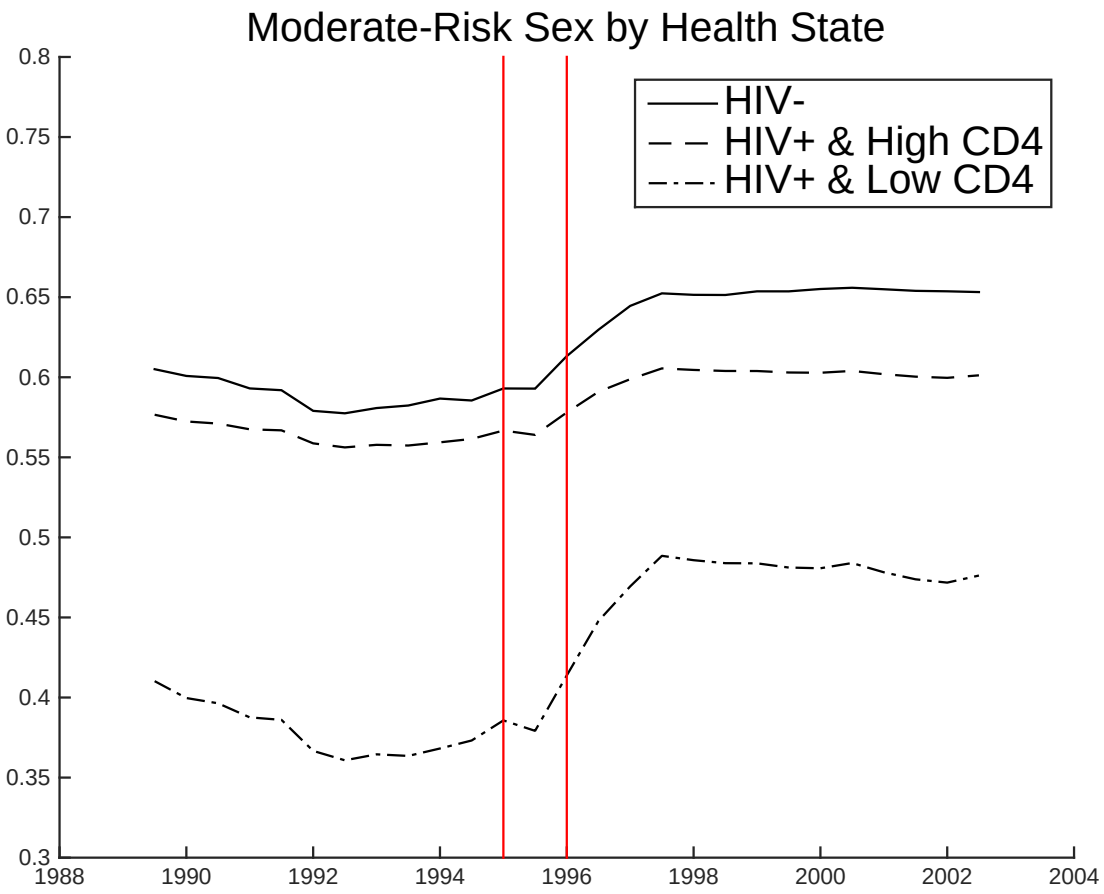

Figure 8: Simulated probability of choosing moderately risky sex behavior over time and for each health state, where the two vertical lines indicate the period of HAART introduction.

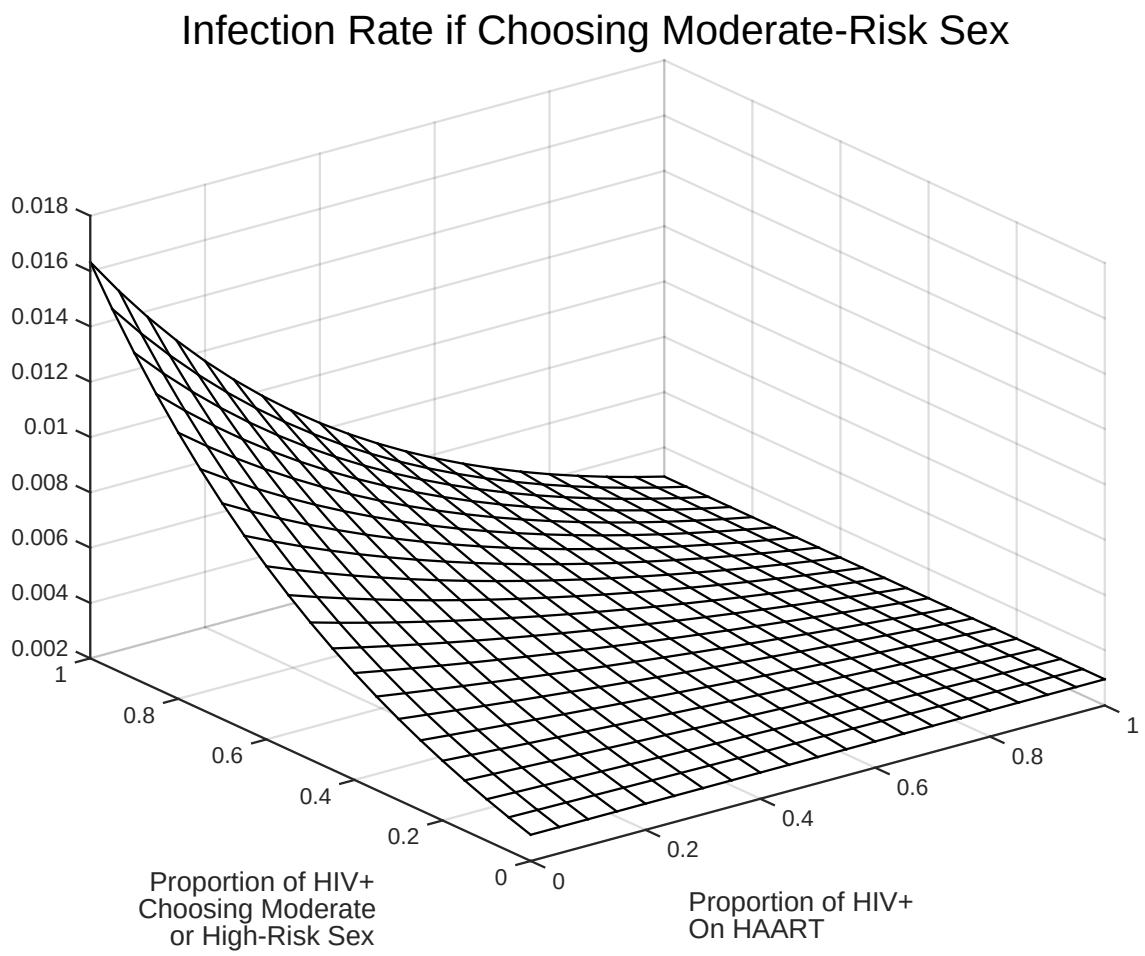

Figure 9: Seroconversion as a function of market-level parameters.

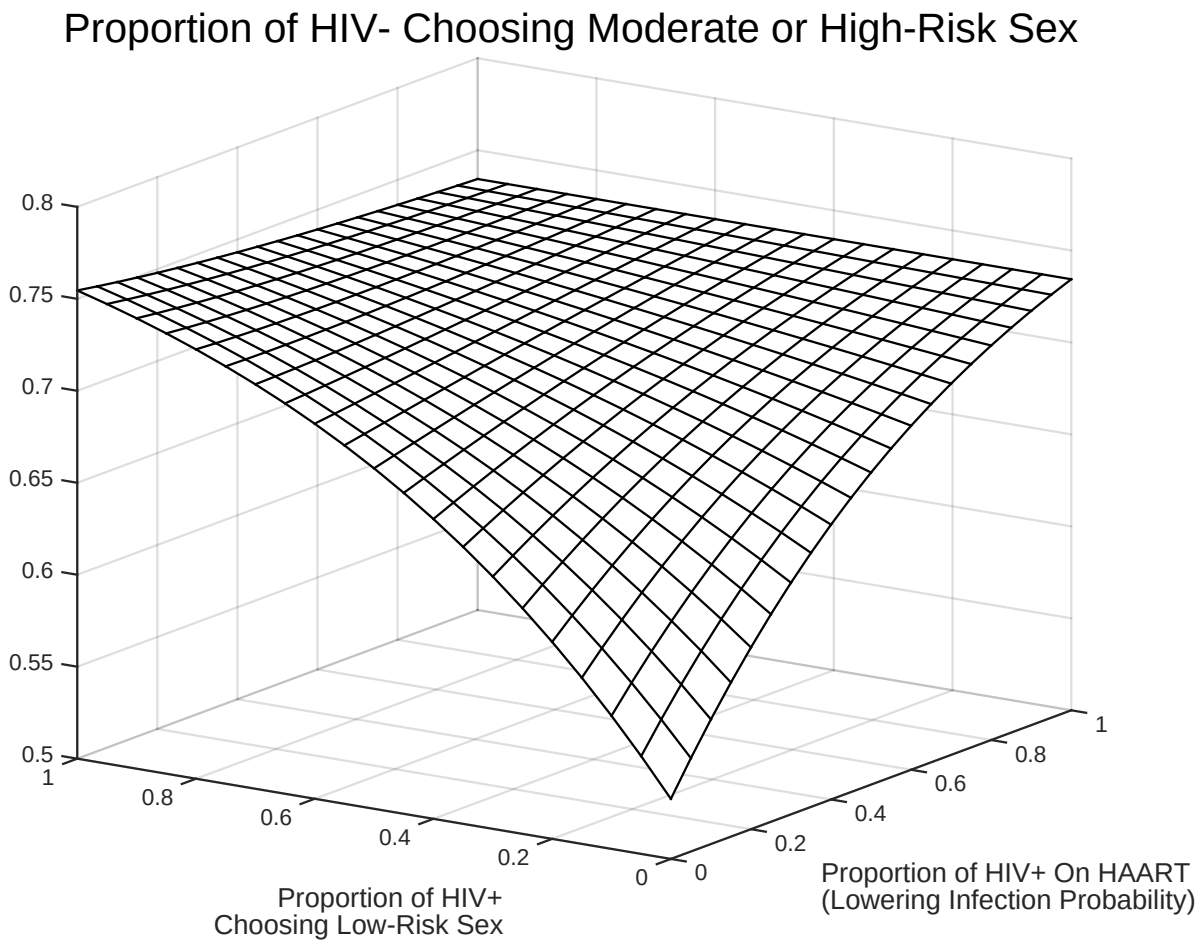

Figure 10: HIV− men's choice to engage in either moderate or high-risk sex as a function of market-level quantities.

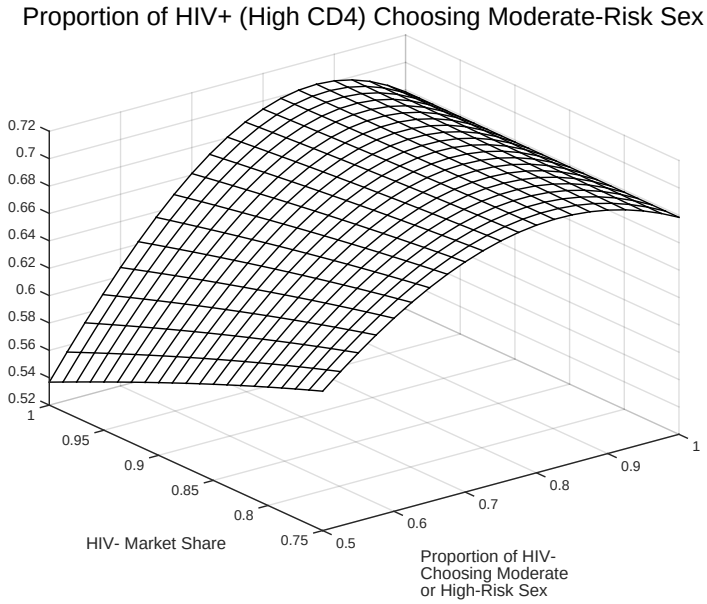

Figure 11: HIV+ sex behavior and the effort function.

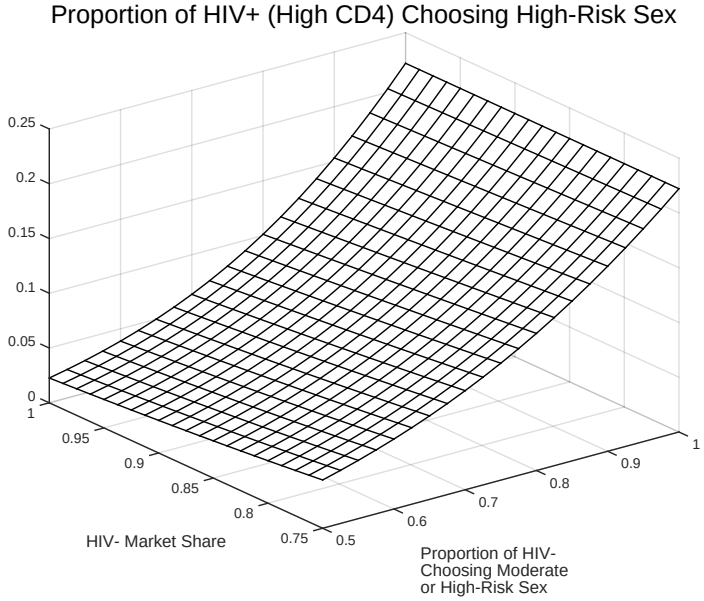

Figure 12: HIV+ sex behavior and the effort function.

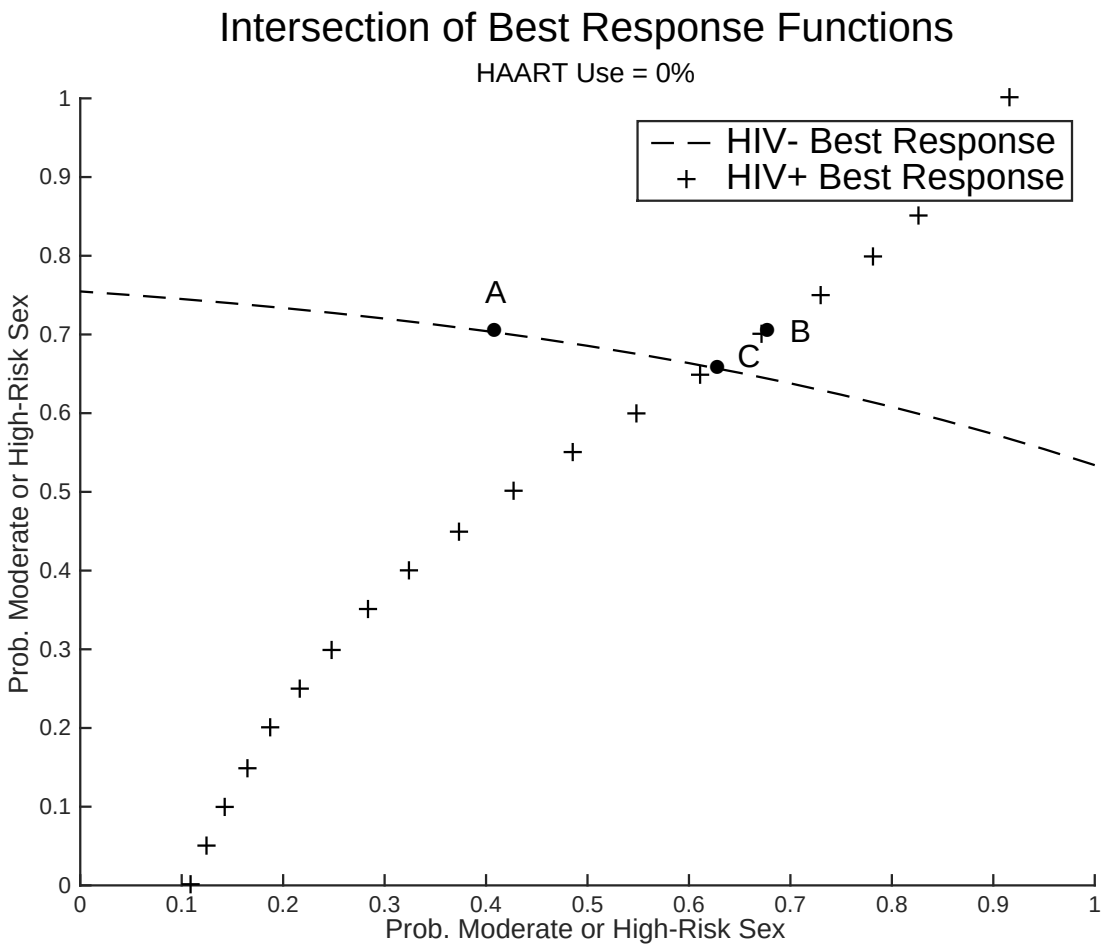

Figure 13: Best responses of HIV− men to market-level HIV+ behavior and best responses of HIV+ men to market-level HIV− behavior. Market-level behavior and best responses are both measured as the probability of moderate or high-risk sex. For HIV− men, the x-axis is market-level HIV+ behavior and the y-axis is the best response. For HIV+ men, the axes are flipped so that market-level HIV− behavior is on the y-axis and best responses of HIV+ men are on the x-axis. Market equilibrium is where best response functions intersect (Point C). At Points A and B, behavior is not at equilibrium.

A Data Appendix

In this section, we describe how we calculate the probability that a potential partner on the MSM partner market is HIV+ with a high or low CD4 count. Quantities used in our computation are summarized in Table A1.[51] First, we calculate the prevalence of HIV in the four cities from which subjects are drawn (Column [1]). For each city, we count the number of HIV cases and divide by the population.[52] Second, we use CDC data to calculate the probability in each year that the cumulative number of AIDS diagnoses consist of MSM (Column [2]). This number is a proxy for the probability that a person living with HIV is a MSM and is used since it is a measure gathered consistently by the CDC over time. An alternative measure would use county or city statistics on the number of people living with HIV who are MSM. Another option is to use CDC data on HIV prevalence, which is provided after 2002. Using alternative measures, the quantity of interest remains quite similar. Third, we approximate the proportion of the U.S. population that can be classified as MSM as 2%. Finally, using Bayes' rule, we calculate:

$$\text{P[HIV+ | MSM]} = \frac{\text{P[MSM | HIV+]} \times \text{P[HIV+]}}{\text{P[MSM]}}, \tag{22}$$

where P[MSM | HIV+] is proxied using P[MSM | AIDS Diagnosis]. This way, we obtain our preferred measure of the likelihood that an MSM is HIV+, which is displayed in Column [4]. In Columns [5]-[6], we present the proportion of HIV+ individuals who have high versus low CD4 counts. These quantities are from the MACS data. Finally, we multiply the quantities in Column [4] with those in Columns [5] and [6] to compute the probability that an MSM is HIV+ with a high CD4 count (Column [7]) or a low CD4 count (Column[8]).

An alternative would be to only use county level statistics for one of the four MACS sites (Baltimore, Chicago, Los Angeles and Pittsburg). County level statistics include the portion of people living with HIV who are MSM. Our preferred measure is an average over the four cities, which means we may over-estimate the probability that subjects in Los Angeles face a positive partner, but under-estimating the same quantity for subjects in Baltimore. We prefer the average because we cannot identify the sites from which individuals in the MACS data set are drawn from. Further, it is not clear which information individuals relied upon

[51] Note that we calculate this quantity for each half-year period, but only present annual numbers to save space.

[52] For Baltimore, we use the Baltimore City HIV/AIDS Epidemiological Profile (Fourth Quarter 2010); for Chicago, we use the 2013 HIV/STI Surveillance report, Chicago; for Pittsburgh, we use the HIV/AIDS Surveillance Biannual Summary for the commonwealth of Pennsylvania; for Los Angeles, we use the 2004 Los Angeles Epidemiological HIV/AIDS Profile. For all cities, we use the U.S. census for total population.

and we assume that national-level statistics were readily available.

Table A1: Computing P[HIV+ | MSM].

Year	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
1989	0.001	0.675	0.020	0.042	0.671	0.329	0.028	0.014
1990	0.001	0.661	0.020	0.046	0.661	0.339	0.031	0.016
1991	0.002	0.648	0.020	0.053	0.650	0.350	0.034	0.018
1992	0.002	0.636	0.020	0.068	0.617	0.383	0.042	0.026
1993	0.002	0.610	0.020	0.070	0.615	0.385	0.043	0.027
1994	0.002	0.591	0.020	0.072	0.562	0.438	0.041	0.032
1995	0.003	0.578	0.020	0.076	0.537	0.463	0.041	0.035
1996	0.003	0.566	0.020	0.078	0.600	0.400	0.047	0.031
1997	0.003	0.553	0.020	0.079	0.662	0.338	0.052	0.027
1998	0.003	0.544	0.020	0.082	0.768	0.232	0.063	0.019
1999	0.003	0.536	0.020	0.085	0.769	0.231	0.065	0.020
2000	0.003	0.528	0.020	0.088	0.757	0.243	0.066	0.021
2001	0.004	0.521	0.020	0.092	0.762	0.238	0.070	0.022
2002	0.004	0.517	0.020	0.099	0.825	0.175	0.082	0.017

This table summarizes how we compute the probability that a partner on the market faced by MSM in a given year is HIV+ and has a high or low CD4 count. To compute this number, we use data from several sources. Column [1]: county level data on HIV prevalence rates, aggregated for the four sites where MACS data are collected. Column [2]: CDC data on the proportion of cumulative AIDS diagnoses that are from individuals who are MSM. This is a proxy for the probability that a person living with HIV is a MSM and is used since it is a measure gathered consistently by the CDC over time. Column [3]: CDC data on the proportion of adults in the U.S. who are MSM, which we assume is constant over the sample period. Column [4]: using data from [1]-[3] and Bayes' rule, we compute the probability that a sex partner on the MSM partner market is HIV+. Columns [5]-[6]: MACS data on the relative weights of HIV+ men who have high versus low CD4 counts. Columns [7]-[8]: Entries in column [4] multiplied by entries in Columns [5]-[6] respectively, giving the probability that a sex partner on the market for MSM is HIV+ with a high or low CD4 count, respectively.

B Model Appendix

In this appendix, we provide specifications for functions that govern survival, HIV+ health state transitions and the decision to use HAART. We note that agents age one half year in each period and remaining state variables (anonymous sex prior to the sample period, college education and race) are fixed. In each period, agents take account of the fact that

they might die between prior to the end of period $t+1$. We model the probability of death as follows. For $h = h_{it} \in \{0, 1, 2\}$,

$$\mathrm{P}(B_{it} = 1|h, A_{it}, C_i, R_i) = \frac{\exp(\theta^{B1} A_{it} + \theta^{B2h} h + \theta^{B3} C_i + \theta^{B4} R_i + \theta^{B5})}{1 + \exp(\theta^{B1} A_{it} + \theta^{B2h} h + \theta^{B3} C_i + \theta^{B4} R_i + \theta^{B5})} \tag{23}$$

HIV+ agents at each period have a healthy (high) CD4 count or AIDS-level (low) CD4 count. They can transition between these two states and the transition probabilities are influenced by HAART use. We model HIV+ health state transitions between the two states as follows. For $j \in \{1, 2\}$ and $h_{i,t+1} \in \{1, 2\}$,

$$\mathrm{P}(h_{i,t+1} \neq j|h_{it} = j, A_{it}, C_i, R_i) = \frac{\exp(\theta_j^{h1} A_{it} + \theta_j^{h2} C_i + \theta_j^{h3} R_i + \theta_j^{h4} H_{it} + \theta_j^{h5})}{1 + \exp(\theta_j^{h1} A_{it} + \theta_j^{h2} C_i + \theta_j^{h3} R_i + \theta_j^{h4} H_{it} + \theta_j^{h5})} \tag{24}$$

Notice that the j index means that separate parameters are estimated for transitions from high to low CD4 counts and from low to high CD4 counts.

We use a reduced-form specification to model the probability that HIV+ individuals use HAART. HIV− individuals, when forming expectations over HAART use in case they ever become HIV+, use the same specification. We model the HAART usage probability as follows, where $h = h_{it} \in \{1, 2\}$:

$$\mathrm{P}(H_{it} = 1|h, A_{it}, C_i, R_i) = \frac{\exp(\lambda_{ht}^H + \theta_h^{H1} A_{it} + \theta_h^{H2} C_i + \theta_h^{H3} N_i + \theta_h^{H4} R_i)}{1 + \exp(\lambda_{ht}^H + \theta_h^{H1} A_{it} + \theta_h^{H2} C_i + \theta_h^{H3} N_i + \theta_h^{H4} R_i))} \tag{25}$$

where λ_{ht}^H is a health-state-specific time dummy variable. In periods prior to HAART introduction, these dummy variables are not estimated.

C Estimation Appendix

In this appendix, we describe the estimation algorithm and computation of standard errors.

C.1 Estimation Algorithm

Recall that parameters to be estimated are collected into a vector denoted Ξ:

$$\Xi \equiv \left\langle \theta^u, \theta^e, \theta^X \right\rangle \tag{26}$$

We estimate Ξ via maximum likelihood and construction of the likelihood occurs in several steps. Some quantities used in the algorithm are directly computed from either the sample or external data in a separate first step since they do not change as model parameters change.

Step 1: Compute quantities that do not change as model parameters change.
These quantities are probability functions computed from our sample and from external data sources.

- $\hat{\mathrm{P}}_t(H|h,X)$ [sample data]
  $\implies \hat{\mathrm{P}}_t(H|h)$
- $\hat{\mathrm{P}}(h'|h,H,X), h,h' \in \{1,2\}$ [sample data]
  $\implies \hat{\mathrm{P}}(h'|h,X), h,h' \in \{1,2\}$
  $\implies \hat{\mathrm{P}}(h'|h), h,h' \in \{1,2\}$
  Note: the latter two quantities are time-varying as they aggregate over HAART use, which is time-varying.
- $\hat{\mathrm{P}}(B=1|h,H,X)$ [sample data]
- $\hat{\mathrm{P}}_t(h'=1|h=0,X,d)$ [sample data]
  Note: this quantity is the time-varying probability of seroconversion conditional on sex choice. It is used as the dependent variable to compute Φ^m (the mapping from market-level quantities like equilibrium sex choices to likelihood of infection) in a way to be explained below.
- $\hat{\mathrm{P}}_t(h)$ [external data]
  Note: this quantity is the distribution of health status in different years among men who have sex with men. Because our data set over-samples HIV+ men, this quantity would be biased if computed from the MACS sample.

Step 2: Estimate the model.
Using the quantities computed in *Step 1*, we estimate model parameters Ξ in an iterative way. Model estimation consists of an outer algorithm and an inner algorithm.

Outer Algorithm:

(1) Given any trial value of Ξ and an initial guess of $\mathrm{P}^0(d|h,H,X)$ (here, assuming that conditional choice probabilities are not time-varying), start the Inner Algorithm, where iteration will be explained below.
Inner Algorithm: At the mth iteration:

(i) Use $\mathrm{P}_t^m(d|h,H,X)$ to calculate $\mathrm{P}_t^m(d|h)$ and $\mathrm{P}_t^m(d)$ for the entire sample and then separately for men who are older than 40 and younger than 40, collecting the latter set of probabilities collected into a matrix denoted $\mathrm{P}_t^{m(Age)}(d|h)$. Note that $\mathrm{P}_t^m(d|h,H,X)$ is an equilibrium probability function that is time-varying since it varies with aggregate market-level quantities ($\bar{h}_t$ and $\bar{H}_t$), which change over time.[53]

(ii) Using a logit model, estimate the parameter Φ^d in the following equation:

$$\mathrm{P}_t(h^{'}=1|h=0,X,d) \;=\; \Phi^d \times \mathrm{P}_t^m(d|h\neq 0)\mathrm{P}_t(H=0|h\neq 0)\mathrm{P}_t(h\neq 0) \quad (27)$$

where estimates are denoted $\hat{\Phi}^{dm}$ and

$$\hat{\Phi}^m = \left\langle \hat{\Phi}^{0m}, \hat{\Phi}^{1m}, \hat{\Phi}^{2m} \right\rangle. \quad (28)$$

Note that $\hat{\Phi}^m$ defines a mapping from market-level quantities to seroconversion probabilities.[54]

(iii) Calculate elements of the effort function $e_j^m(\hat{\mathrm{P}}_t(h=0), \mathrm{P}_t^m(d=0|h=0), \Xi)$.

(iv) Using Φ^m and equation (27), calculate seroconversion probability from period t and $t+1$

$$\hat{\mathrm{P}}_t^m(h'_{it}=1|h_{it}=0,X_{it},d_{it}) \quad (29)$$

(v) Next, use equation (16) to calculate the mth iteration of $\bar{V}_j(Z)$, which we denote $\bar{V}_j^m(Z)$.

(vi) Use $\bar{V}_j^m(Z)$ to calculate $\mathrm{P}_t^{m+1}(d=j|h,H,X)$ according to equation (17), where $\mathrm{P}_t^{m+1}(d=j|h,H,X)$ is the $(m+1)$-th iteration of $\mathrm{P}_t(d=j|Z)=\phi_j(Z)$ from equation (17).

(vii) Compute $||\mathrm{P}_t^{m+1}(d|h,H,X) - \mathrm{P}_t^m(d|h,H,X)||$.

- If $||\mathrm{P}_t^{m+1}(d|h,H,X) - \mathrm{P}_t^m(d|h,H,X)|| \leq \epsilon_p$, then stop. This implies that, given Ξ, the equilibrium constraints are satisfied because beliefs about $\mathrm{P}(\tilde{d}|\tilde{h}) = \mathrm{P}(d|h)$ used in calculating the value function for an agent when making sex choices is consistent with the equilibrium sex choice function $\mathrm{P}_t^{m+1}(d|h,H,X)$. Go to the outer algorithm (step (2) below);

[53] In general, the equilibrium choice probability function also varies with changes in aggregate X values since it is a function of aggregate sex probabilities, which are affected by X.

[54] Equation (27) comes from rewriting equation (15) imposing that unobserved partners exhibit the same choice and transition probabilities as the individuals in our sample, as discussed when we listed equilibrium constraints. Further, note that in the implementation of the model, we also estimate an intercept term in the model described by equation (27).

- If $||\mathrm{P}_t^{m+1}(d|h,H,X) - \mathrm{P}_t^m(d|h,H,X)|| > \epsilon_p$, then, replace $\mathrm{P}_t^m(d|h,X)$ by $\mathrm{P}_t^{m+1}(d|h,X)$ and then start iteration $m+1$ by repeating procedure (ii) to (vii) until convergence.[55]

(2) Evaluate the likelihood function. Given our modeling assumptions, we only need to match the sex choice probabilities. The likelihood contribution for individual i at time t is given by:

$$L_{it}^P(Z_{it},\Xi) = \prod_d \phi_d(Z_{it},\Xi)^{\mathbf{1}[d_{it}=d]} \tag{30}$$

Here, $\phi_d(Z_{it},\Xi)$ is the probability of choice d implied by the model given parameters Ξ and state variables Z_{it} and $\mathbf{1}[\cdot]$ is the indicator function for individual i at time t actually choosing d. Since market variables are time-varying, we match choice probabilities implied by the data for each period to analogous period-specific moments in the data. Repeat the inner procedures for different trial values of Ξ until the likelihood is maximized.

C.2 Computation of Standard Errors

We compute standard errors by constructing the Hessian of the likelihood function using the outer product measure. To compute the outer product measure, we calculate two-sided numerical derivatives of the likelihood function for each estimated parameter. In each direction, the derivative is calculated by perturbing each parameter and then solving for equilibrium choice probabilities under the equilibrium constraint implied by the perturbed parameter and then computing the likelihood. Notice, this method of calculating numerical derivatives takes full account of how estimated parameters affect not only optimal responses under equilibrium constraints, but also the constraints themselves. Finally, although the estimation algorithm proceeds in two stages, where first-stage estimates are taken as fixed in solving for second stage parameters, we compute errors using the full likelihood function that includes both first and second stage parameter estimates. In this way, we take full account of how first-stage estimates are measured with error.[56]

[55] Note that there is no guarantee that this procedure will converge given Ξ. The empirical procedure must be tested with different parameter values during model estimation.

[56] Since county-level data represent the population of agents with HIV, we take these quantities as fixed.

D Decomposition Appendix

In this appendix, we describe how we isolate the proportion of the value of HAART that is attributable solely to its impact on survival conditional on HIV infection. To understand how we do this, first note that the effect of HAART on HIV+ men is to lower the amount of time they spend with AIDS. In other words, HAART does not make AIDS less deadly, but less likely. This has two direct effects. First, it raises the likelihood of survival conditional on HIV infection. Second, it improves quality of life in terms of higher flow utility and lower effort function parameters for the incremental time spent in the non-AIDS state. There are also market equilibrium effects. Higher survival probability and better quality of life (as measured by higher utility) lower the cost of HIV infection from the perspective of HIV− agents, who respond by increasing their risky sex behavior. This benefits HIV+ men who face lower effort costs of finding willing sexual partners. Here, our aim is to focus solely on the value of HAART from its effect on survival, which means we subtract the value attributable to higher quality of life and from market equilibrium effects. Isolating the value of HAART arising from survival proceeds according to the following steps. We do not in any step allow agents to re-optimize or allow market-level quantities to change from their pre-HAART levels.

(1) Compute the total continuation payoff of HAART under the constraint that agents do not re-optimize from pre-HAART behavior and given pre-HAART market quantities. We obtain a value of HAART that only includes its impact on survival and higher period utility for HIV+ men.

(2) Repeat the calculation in (1), but impose that in both HIV+ health states agents face AIDS-level utility and effort function parameters. This calculation delivers a value of HAART that is too low. The reason is that we only want to subtract gains from higher utility from HAART-induced *incremental* time spent in the non-AIDS state. In this step, we have assigned AIDS-level utility to *all* time spent in the non-AIDS state.

(3) Calculate pre-HAART value.

(4) Calculate value using pre-HAART health state transition probabilities, but imposing that both HIV+ health states deliver AIDS-level CD4 count utility.

(5) Subtract the value calculated in step (4) from the value calculated in step (3). The result tells us how much value is lost by imposing AIDS-level utility parameters onto the number of pre-HAART periods HIV+ agents expect to be HIV+ and in the non-AIDS state.

(6) Subtract the value computed in step (5) from the value computed in step (2). The result tells us how much utility is lost by imposing AIDS-level utility for the HAART-induced incremental expected time spent in the non-AIDS state.

(7) Subtract the value that is computed in step (6) from the value calculated in (1). What is left is the value of HAART that is solely attributable to survival and that is not due to higher utility parameters, re-optimization and market equilibrium effects.

This leaves us with the value of HAART due to its effect on survival conditional on HIV-infection. For HIV− men, that means that 18% of the total value of HAART is from its effect on survival conditional on HIV-infection. For HIV+ individual with high or low CD4 counts, 77% of the value of HAART is attributable to the effect on survival.