The Marriage Market for Lemons: HIV Testing and Marriage in Rural Malawi

Manuela Angelucci*
Daniel Bennett†
Jenny Trinitapoli‡

February 29, 2016

Abstract

In HIV-endemic settings, HIV status is an important hidden attribute in the marriage market. Asymmetric HIV status information may cause adverse selection by discouraging marriage among healthy people. This paper develops a simple assortative matching model with adverse selection on health. We test the model through a randomized evaluation of an intensive opt-out HIV testing intervention, which offered tests to young women and their partners every four months for 2.5 years. Consistent with the model, the intervention significantly increased marriage and pregnancy, as well as awareness of partner health. We show heterogeneous effects by HIV status and physical attractiveness (another key marriage market attribute) that match our theoretical predictions. Finally, we relate our findings to the literature by showing that an alternative single-test intervention does not have these effects.

JEL: J12, J13, I15, I18

We received helpful feedback from Dan Black, Alessandra Voena, and Sara Yeatman. Hannah Furnas, Ashley Larsen Gibby, and Ngoc Nguyen provided excellent research assistance. This research uses data from the Tsogolo la Thanzi Panel Study, funded by grants (R01-HD058366 and R01-HD077873) from the National Institute of Child Health and Human Development.

*Department of Economics, University of Michigan, mangaluc@umich.edu.
†Harris School of Public Policy Studies, University of Chicago, dmbennett@uchicago.edu
‡Department of Sociology, University of Chicago, jennytrini@uchicago.edu.

1 Introduction

HIV risk is an important partner trait for marriage market participants in HIV-endemic settings. HIV risk encompasses both current HIV status and future infection risk. A low-risk spouse is desirable because he or she contributes to both economic and financial security, and reduces own HIV risk. However partner risk profiles are difficult to observe. HIV risk indicators such as AIDS symptoms, the use of antiretroviral therapy, and non-monogamous behavior are only observable with time and repeated interaction. As a result, low-risk people must either postpone marriage or choose a partner from an adversely-selected pool.

Signaling and screening, the textbook ways to resolve adverse selection, are not generally available in terms of HIV because because HIV testing remains costly. While testing is nominally free in much of sub-Saharan Africa, access is limited in practice by long distances to clinics and long delays in returning test results (Matovu and Makumbi 2007). Pinto et al. (2013) report that HIV-related visits take an average of 7.1 hours in rural Malawi due to lengthy travel and wait times. Seeking an HIV test is also stigmatizing because it suggests to observers that the test recipient may have been promiscuous (Chesney and Smith 1999, Young and Zhu 2012). The costliness of HIV testing encourages a pooling equilibrium in which healthy people cannot use HIV test results to signal and screen.

A reduction in the cost of testing may influence the marriage market by resolving the information asymmetry. In an environment with costless testing, low-risk people can signal to partners by sharing test results and can screen partners by the willingness to be tested. To facilitate signaling and screening, the cost reduction must be dramatic and sustained. Many HIV testing interventions in the literature offer HIV tests on a one-off basis (Thornton 2008, Delavande and Kohler 2012, Baird et al. 2014, Beegle et al. 2015). The provision of HIV testing in this way is unlikely enable signaling and screening because it does not change the perceived availability of HIV status information for marriage market participants going forward.

To begin, we develop a simple model of assortative matching under asymmetric information. We assume that people know their own HIV status but cannot initially observe the HIV status of potential partners. In the status quo, HIV positive people marry immediately but some HIV negative people postpone marriage, leading to adverse selection in the marriage market in Period 1. People in the model have another attribute, attractiveness, which is always observable. A complementarity between attractiveness and health (or higher HIV prevalence among attractive people) leads to the prediction that attractive people delay marriage the most. A change that provides information about the HIV status of potential partners (such as the intervention we describe below) resolves the information asymmetry and accelerates marriage for HIV-negative people. This intervention increases the marital surplus of healthy people. However we show that, insofar as people assortatively match on attractiveness in the status quo, the intervention does not cause HIV-negative people to obtain more attractive partners.

We test this model by evaluating the marriage market impact of an intensive HIV-testing intervention in rural Malawi. The Tsogolo La Thanzi (TLT) Panel Study follows a representative sample of 1500 young women across eight waves from 2009 to 2012. Participants in the randomly-assigned treatment arm were offered a free HIV test at the end of every survey wave. Surveyors urged participants to recruit their partners into the study, and partners of treatment participants were also offered HIV tests. We argue that this intervention helped to resolve the information asymmetry related to partner health. In contrast to existing HIV testing interventions, which typically offer tests only once to one partner, this intervention reduced the expected cost of HIV testing for both partners. Using detailed partnership records and pregnancy biomarkers, we show that this intervention increased marriage by 3.5 percentage points and pregnancy incidence by 2.1 percentage points.

Several additional tests allow us to assess the model further. While the intervention has a minimal effect on own perceived HIV status, it improves the precision of partner HIV status perceptions. The model also predicts several patterns, which we test using baseline data. Under some assumptions, fertility serves as a proxy for marital surplus. We show that HIV-negative and attractive women have higher fertility per year of marriage. We use baseline surveyor observations of respondent physical attractiveness to examine the possible interaction between observable and unobservable marriage market attributes. Consistent with a complementarity between these attributes, the intervention has the largest impact on attractive, HIV-negative women. For this group, the intervention increases the probability of marriage by 10 percentage points and the probability or pregnancy by 4.5 percentage points.

Finally, we contrast our findings with existing HIV testing evaluations. Studies generally offer a single HIV test and assess impacts on risky sexual behavior, and to a lesser extent, marriage, education, and fertility. Impacts are small and appear to depend on the test result and whether the recipient was surprised by the news.[1] Our study includes a third experimental arm that was offered HIV testing on only one occasion. We find no impacts of this intervention on marriage or fertility, suggesting that the repeated testing is an important feature of the intervention.

This paper makes the following primary contributions. We provide one of the first empirical applications of adverse selection in the marriage market. Unobservable partner quality is an intrinsic feature of the marriage market in many settings. For instance, female virginity is highly valued in many conservative settings but is difficult to observe. The inability to signal this attribute may influence marriage timing and schooling decisions. Partner finances are also difficult to observe before marriage. We suggest that unobservable HIV risk leads healthy people to delay marriage and fertility. Secondly we contribute to the literature on HIV testing by showing a novel impact of routine testing. Policymakers have debated how to offer HIV testing in sub-Saharan Africa. Analyses have focused on the way that test results may inform individual choices rather than their possible role as signals to others. The impact on the marriage market merits further policy consideration.

[1] Thornton (2008) shows that HIV testing modestly increases condom demand. Baird et al. (2014) find that testing negative in a home-based intervention does not change the prevalence of sexually-transmitted infections (STIs) but that testing positive increases STI prevalence. Gong (2015) shows that positive test results increase STI infections and negative results decrease STI infections, but only for people who are surprised by the results. Beegle et al. (2015) find no effect of a one-off testing intervention on school attendance, marriage, and fertility.

2 Theory

2.1 Setup

Consider a transferable utility setting with a continuum of men and women who live for two periods, $t \in \{1, 2\}$. In each period, people decide whether to participate in the marriage market and whom to marry. People receive surplus S for one period of marriage and face discount factor $\delta > 0$, which can be greater than one because Period 2 lasts longer than Period 1. The discount factor is distributed uniformly between 0 and b: $\delta \sim U[0, b]$. Therefore, a person who marries in Period 1 has lifetime surplus $(1 + \delta)S$. People do not die and divorce is too costly, so that people who marry in Period 1 remain with their partners in Period 2.

People have two fixed binary traits, attractiveness and health, which may be either high or low (h or l). Therefore, there are four types of agents, defined by their attractiveness and health, with population shares p_{hh}, p_{lh}, p_{hl}, and p_{ll}, that are common knowledge. Attractiveness is observed by everyone in both periods. This dimension represents any other salient partner attribute, such as physical beauty for women and wealth for men. While people always know their own health status, the health status of others becomes observable in Period 2. If a woman with attractiveness a and health b and a man with attractiveness c and health d marry, they generate a surplus $S^{ab}_{cd} > 0$ (where the superscript references the woman and the subscript references the man).

Attractive and healthy spouses generate more surplus than, respectively, unattractive and unhealthy ones. From the woman's perspective, a woman of attractiveness a and health b would enjoy the highest surplus by marrying a man with two high traits, the second highest surplus by marrying a man with one high trait, and the lowest surplus by marrying a man with no high traits, $S^{ab}_{hh} > S^{ab}_{hl} = S^{ab}_{lh} > S^{ab}_{ll} > 0$. Finally, in the population, attractiveness, health, and the discount factor are independent of gender, and attractiveness and health are independent of the discount factor.

2.2 Status Quo Equilibrium

Other people's health is unobservable in Period 1. Therefore, people must decide whether to marry in Period 1 or postpone marriage until Period 2 according to their beliefs of the distribution of healthy and unhealthy people in the pool at this point. People have correct beliefs in equilibrium, that is, they correctly perceive the shares of each type who seek to marry in each period. The remainder of this section considers women's choices, to be consistent with our empirical exercise.

This setup generates a unique equilibrium in the marriage market in each period, in which people marry assortatively in attractiveness in Period 1 and in the number of high traits in period 2, all unhealthy women marry in Period 1 and some healthy women marry in Period 2, inducing adverse selection in Period 1. We discuss the details in the following five Propositions.

Proposition 1. *For all women who choose to marry in Period 2, there exists a unique and stable equilibrium in Period 2 in which people marry assortatively in the number of high traits.*

Proof. Since attractiveness, health, and the discount factor have the same the distribution for men and women and are independent of each other, there are equal numbers of marriageable men and women with the same attractiveness and health in each period. Therefore, attractive and healthy people marry each other and enjoy a surplus of S^{hh}_{hh}. Similarly, women with one high trait also marry men with one high trait and enjoy a surplus of S^{lh}_{hl}.[2] Therefore, people without high traits must also marry each other and enjoy a surplus of S^{ll}_{ll}. This equilibrium is stable, that is, no married person are better off unmarried, and no two people prefer being married together than to their current spouse. This occurs because, while any person would prefer to marry a partner with more high traits than one's own, nobody wants to marry a partner with fewer high traits than one's own.

[2]Note that all the following surpluses are identical by assumption: $S^{lh}_{hl} = S^{lh}_{lh} = S^{hl}_{hl} = S^{hl}_{lh}$.

Proposition 2. *There exists a unique and stable equilibrium in Period 1 in which people marry assortatively in attractiveness.*

Proof. For women who choose to marry in Period 1, attractive women maximize surplus by marrying attractive men because the surplus from marrying an attractive partner is always (weakly) greater than the surplus from marrying an unattractive partner, since the former has (weakly) more high traits in expectation. Given that attractive people marry each other, unattractive people must marry each other too.[3] If divorce costs were sufficiently low, this equilibrium would be unstable, since all people married to partners with fewer high traits than their own would prefer being married together to their current situation. We discuss this in section 6.

Proposition 3. *Unhealthy women always marry in Period 1.*

Proof. Unhealthy women of attractiveness a, al types, marry early if the surplus from doing so is greater than the surplus from marrying late, that is, if the following inequality holds:

$$(1+\delta)\frac{p_{al}S_{al}^{al}+p_{ah}S_{ah}^{al}}{p_{al}+p_{ah}} > \delta S_{al}^{al}, \tag{1}$$

where the variables p_{al} and p_{ah} are the population proportions of low and high health types with attractiveness a. Since $S_{ah}^{al} > S_{al}^{al}$, this expression is always true. Unhealthy women who marry in Period 1 receive additional surplus in two ways. First, they are married for longer, so they generate marital surplus for more time. Second, by marrying early, they may find a healthy spouse, who is unaware that they are unhealthy, and, obtain a higher marital surplus than if they married an unhealthy partner. This pattern does not depend on attractiveness. That is, both attractive and unattractive unhealthy women marry in Period 1.

[3] Marriage may not be assortative in attractiveness if health and attractiveness are correlated. For example, if they are negatively correlated, attractive people may prefer to marry unattractive people if their expected health is sufficiently higher. In this scenario, unlike in our model, HIV testing would cause people to match assortatively in attractiveness.

Proposition 4. *Healthy women marry in Period 1 if their discount factor is low enough.*

Proof. The population proportion of healthy women of attractiveness a who marry in Period 1, μ_{ah}, can vary between 0 and the entire population proportion p_{ah}, $\mu_{ah} \in [0; p_{ah}]$. The parameter μ_{ah} depends on the marriage surpluses and the discount factor δ. We derive $\bar{\delta}_a$ considering the choice of healthy women of attractiveness a, who marry early if the expected surplus of marrying in Period 1 exceeds the surplus of marrying in Period 2, that is:

$$(1+\delta)\frac{p_{al}S_{al}^{ah} + \mu_{ah}S_{ah}^{ah}}{(p_{al} + \mu_{ah})} > \delta S_{ah}^{ah}. \tag{2}$$

Solving for δ yields an expression for $\bar{\delta}_a$, the threshold value that distinguishes healthy women who marry early and late. Healthy women of attractiveness a marry in Period 1 if they are sufficiently impatient that $\delta < \bar{\delta}_a$, where:

$$\bar{\delta}_a = \frac{\mu_{ah}S_{ah}^{ah} + p_{al}S_{al}^{ah}}{p_{al}(S_{ah}^{ah} - S_{al}^{ah})} > 0 \tag{3}$$

Equation (3) is always positive, because both its numerator and its denominator are positive, and depends on μ_{ah}, the share of other healthy people who marry early.

Since δ is distributed uniformly between 0 and b, we know that $\mu_{ah} = F(\bar{\delta}) = \bar{\delta}_a/b$ or that $\mu_{ah}b = \bar{\delta}_a$. Setting this expression equal to Equation (3) and solving for μ_{ah} gives:[4,5]

$$\mu_{ah}^{*} = \frac{p_{al}S_{al}^{ah}}{bp_{al}(S_{ah}^{ah} - S_{al}^{ah}) - S_{ah}^{ah}} \tag{4}$$

Proposition 5. *Under some conditions, a higher proportion of attractive and healthy women, types hh, marry in Period 2, than unattractive and healthy women, types lh. That is,* $\mu_{hh} < \mu_{lh}$.

[4]Recall that $0 \leq \mu_{ah}^{*} \leq p_{ah}$. Therefore, $\mu_{ah} > 0$ if $b > S_{ah}^{ah}/p_{al}(S_{ah}^{ah} - S_{al}^{ah})$ and $\mu_{ah} \leq p_{ah}$ if $b \geq (p_{al}S_{ah}^{ah} + S_{al}^{ah})/p_{ah}p_{al}(S_{ah}^{ah} - S_{al}^{ah})$.

[5]You can substitute μ_{ah}^{*} into Equation (3) to obtain an expression for $\bar{\delta}_a^{*}$ in terms of b, S_{ah}^{ah} and S_{al}^{ah}: $\bar{\delta}^{*} = \frac{S_{ah}^{ah}S_{al}^{ah}}{bp_{al}(S_{ah}^{ah} - S_{al}^{ah})^2 - S_{ah}^{ah}(S_{ah}^{ah} - S_{al}^{ah})} + \frac{S_{al}^{ah}}{S_{ah}^{ah} - S_{al}^{ah}}$.

Proof. From Equation 4, it follows that $\mu_{hh} < \mu_{lh}$ if:

$$\frac{S_{hh}^{hh}}{S_{hl}^{hh}}(b - \frac{1}{p_{hl}}) > \frac{S_{lh}^{lh}}{S_{ll}^{lh}}(b - \frac{1}{p_{ll}}) \tag{5}$$

This inequality may hold for either of two reasons. First, when $p_{hl} > p_{ll}$ and $\frac{S_{hh}^{hh}}{S_{hl}^{hh}} = \frac{S_{lh}^{lh}}{S_{ll}^{lh}}$. In this case there are more "lemons" among attractive types, i.e., $p_{hl} > p_{ll}$, and the relative gain that healthy women experience by marrying a healthy man rather than an unhealthy man does not vary with attractiveness, i.e., $\frac{S_{hh}^{hh}}{S_{hl}^{hh}} = \frac{S_{lh}^{lh}}{S_{ll}^{lh}}$.[6] Second, when $p_{hl} = p_{ll}$ and $\frac{S_{hh}^{hh}}{S_{hl}^{hh}} > \frac{S_{lh}^{lh}}{S_{ll}^{lh}}$. In this case, the proportion of "lemons" does not vary by attractiveness, but the relative gain that healthy women experience by marrying a healthy man rather than an unhealthy man is larger for attractive women than for unattractive women. This happens if, for example, the current low health of the husband causes a decrease in his future productivity. In that case, the future income generated by the husband would be much lower than the current one.[7]

In our data, $p_{hl} = 0.056$ and $p_{ll} = 0.042$, that is, $p_{hl} > p_{ll}$ but the two probabilites are similar. Therefore, a higher proportion of attractive women prefers to wait to in Period 2 when $\frac{S_{hh}^{hh}}{S_{hl}^{hh}} = \frac{S_{lh}^{lh}}{S_{ll}^{lh}}$, but also in some cases in which $\frac{S_{hh}^{hh}}{S_{hl}^{hh}} > \frac{S_{lh}^{lh}}{S_{ll}^{lh}}$

This equilibrium has the following implications for Period 1, which we can check in our data to validate our theoretical setup:

1. By assumption, marital surplus is bigger for attractive than for unattractive women and for healthy than unhealthy women.

2. According to Proposition 2, the likelihood of being married to an attractive man is higher for attractive women than for unattractive women. Conversely, women's health and husbands' education should not be correlated because HIV status is hard to observe.

[6]This ratio is constant for surplus functions such as $S_{cd}^{ab} = abcd$ or $S_{cd}^{ab} = (a+c)(b+d)$, $\forall l, h > 1$.
[7]This case can be can be illustrated by the following surplus, $S_{cd}^{ab} = (ac)^{bd}$, $\forall l, h > 1$.

3. According to Propositions 3 and 4, the likelihood of being married is higher for unhealthy women than healthy women.

4. According to Proposition 5, the likelihood of being married is higher for unattractive women than attractive women.

2.3 The Impact of Making Health Observable in Period 1

In this Section, we describe how making health observable in Period 1 changes the timing and the surplus of marriages for some women.

First, consider the status quo in which health is unobserved in period 1. In this setting, unhealthy women, al types, marry in Period 1 with probability 1. Their Period 1 expected surplus, $E[S^{al}]_1$ is:

$$E[S^{al}]_1 = \frac{(\mu^*_{ah} S^{al}_{ah} + p_{al} S^{al}_{al})}{p_{al} + \mu^*_{ah}}. \tag{6}$$

Healthy women, ah types, marry in Period 1 with probability $\frac{\mu^*_{ah}}{p_{ah}}$ and in Period 2 with probability $\frac{p_{ah} - \mu^*_{ah}}{p_{ah}}$. Their Period 1 expected surplus, $E[S^{ah}]_1$, is:[8]

$$E[S^{ah}]_1 = \left(\frac{\mu^*_{ah}}{p_{ah}}\right) \frac{(\mu^*_{ah} S^{ah}_{ah} + p_{al} S^{ah}_{al})}{(p_{al} + \mu^*_{ah})}. \tag{7}$$

Suppose that health becomes observable in Period 1. In this case, all women marry a partner with the same number of high traits in Period 1 and nobody postpones marriage to Period 2.

Consider some treatment, T, that makes health observable in Period 1 and some outcome of interest, Y (marital status and surplus, in our case). We can measure the average causal effect of making health observable in Period 1 by comparing the average outcomes when

[8]Lifetime marriage surpluses for unhealthy and healthy women are $E[S^{al}] = \frac{(1+\delta)(\mu^*_{ah} S^{al}_{ah} + p_{al} S^{al}_{al})}{p_{al} + \mu^*_{ah}}$ and $E[S^{ah}] = (\frac{\mu^*_{ah}}{p_{ah}}) \frac{(1+\delta)(\mu^*_{ah} S^{ah}_{ah} + p_{al} S^{ah}_{al})}{(p_{al} + \mu^*_{ah})} + (\frac{p_{ah} - \mu^*_{ah}}{p_{ah}}) \delta S^{ah}_{ah}$.

health is and is not observable. More precisely, define the average causal effect of the treatment on treated women, i.e., on women for whom own and potential partner's health becomes observable in Period 1, with attractiveness a and health b, as the difference in average outcomes in the presence ($T = 1$) and in the absence ($T = 0$) of the treatment, i.e., $ATT(Y^{ab}) = E[Y^{ab}|T = 1] - E[Y^{ab}|T = 0]$.

The treatment has the following average causal effects on marriage, surplus, and matching in Period 1:

The treatment has no effect on marital status for unhealthy women, who always marry in Period 1. That is,

$$ATT(M_1^{al}) = 0. \tag{8}$$

Conversely, the treatment increases the likelihood that healthy women marry in period one ($M_1^{ah} = 1$) from $\frac{\mu^*_{ah}}{p_{ah}}$ to 1. That is,

$$ATT(M_1^{ah}) = 1 - \frac{\mu^*_{ah}}{p_{ah}} \geq 0. \tag{9}$$

This occurs because the women who would have been better off by waiting to find out the health of potential partners can do so in Period 1, and, therefore, marry earlier.[9]

Recall that, in our data, the share of the population which is either attractive and healthy or unattractive and healthy are approximately the same, $p_{hh} = 0.447 \approx p_{lh} = 0.454$. Therefore, the treatment effect is bigger for attractive healthy women than for unattractive healthy women, as long as $\mu^*_{hh} < \mu^*_{lh}$, consistent with our data.

The treatment reduces average Period 1 marital surplus for unhealthy women from $\frac{(\mu^*_{ah}S^{al}_{ah}+p_{al}S^{al}_{al})}{p_{al}+\mu^*_{ah}}$ to S^{al}_{al}. That is,

$$ATT(S_1^{al}) = -\frac{\mu^*_{ah}(S^{al}_{ah} - S^{al}_{al})}{(p_{al} + \mu^*_{ah})} < 0. \tag{10}$$

[9]The treatment has no effect on lifetime marriage likelihood because all women marry by Period 2.

This negative treatment effect occurs because the treatment reveals these women's low health status, and therefore these women can marry unhealthy partners only. Note that the overall average treatment effect on Period 1 surplus is likely positive because there are many more healthy than unhealthy types in the population.

Conversely, the treatment increases the average Period 1 marital surplus (S) for healthy women from $(\frac{\mu^*_{ah}}{p_{ah}})\frac{(\mu^*_{ah}S^{ah}_{ah}+p_{al}S^{ah}_{al})}{(p_{al}+\mu^*_{ah})}$ to S^{ah}_{ah}. That is,

$$ATT(S)^{ah}_1 = \frac{\mu^*_{ah}S^{ah}_{ah}(p_{ah}-\mu^*_{ah})+p_{al}(p_{ah}S^{ah}_{ah}-\mu^*_{ah}S^{ah}_{al})}{p_{ah}(p_{al}+\mu^*_{ah})} > 0. \tag{11}$$

This is the combination of two positive effects. First, healthy women now marry earlier in expectation. Specifically, the fraction $1-\frac{\mu^*_{ah}}{p_{ah}}$ of women now marries in Period 1, while, in the absence of the treatment, would have married in Period 2. For these women, marital surplus increases from zero to S^{ah}_{ah} in Period 1. That is, these women marry the same type of man regardless of the treatment, but do so earlier. Second, the remaining women now all match with healthy men. For women who would have married an unhealthy man in Period 1, the treatment increases their surplus from S^{ah}_{al} to S^{ah}_{ah}. That is, these women marry at the same time regardless of the treatment, but they marry higher health partners.[10]

We propose to consider number of children as a proxy for surplus. If happier couples have more children, then an increase in surplus will cause higher fertility. An alternative explanation for why fertility is a proxy for surplus occurs if investing in children requires upfront (monetary and non-monetary) costs. Therefore, the higher the marriage surplus, the higher the likelihood of being able to pay the upfront cost. If that is the case, then we can indirectly measure the ATT effects described in equations (10) and (11) by estimating the treatment effects on fertility.

We can test the model by observing in our data whether making health observable changes marriage propensity and marital surplus for unhealthy and healthy women in the

[10]The treatment's effect on lifetime surplus is negative for unhealthy women and and positive for healthy ones.

way described above. Since our model studies the process of getting married, we test these hypotheses by looking at the behavior of unmarried women. However, we observe both unmarried and married women in our data. Therefore, it is useful to understand whether making health observable in Period 1 also changes the marital status and surplus for women already married.

When we rule out divorce, the model predicts no treatment effects on marriage, since women are already married. The treatment also has no average effect on surplus for women married at baseline, regardless of whether surplus depends on partner's health (i.e., whether one's partner is healthy) or on knowledge of a partner's health (i.e., whether spouses think their partners are healthy).

If surplus depends on health, then providing knowledge on partner's health has no effect on surplus. Conversely, if surplus varies with knowledge of partner's health, making health observable increases the surplus of some couples and decrease the surplus of others without changing the average surplus of a type. For example, consider married lh women in Period 1 when health is unobserved. Each woman has an expected surplus of $\frac{p_{lh}S^{lh}_{lh}+p_{ll}S^{lh}_{ll}}{p_{lh}+p_{ll}}$. When health becomes observed, women married to lh and ll men find out that their surplus is S^{lh}_{lh} and S^{lh}_{ll}. However, the expected surplus is still $\frac{p_{lh}S^{lh}_{lh}+p_{ll}S^{lh}_{ll}}{p_{lh}+p_{ll}}$.

Lastly, our model makes specific assumptions about risk preferences, the relative values of traits, divorce costs, and the correlation between the traits. With minor exception, making health a larger determinant of surplus than attractiveness, allowing for divorce, letting attractive people be more patient, and letting trait frequencies vary by gender generally does not change our equilibrium and predictions. Appendix 6 considers these cases in more details.

3 Context and Data

3.1 HIV Testing

The cost of HIV testing has gradually declined in sub-Saharan Africa in recent years. With support from international donors, many national ministries of health offer free HIV tests at primary care clinics. While nominally free, this form of testing incorporates substantial transaction costs. Pinto et al. (2013) note that patients in the Zomba District of Malawi, which is nearby the study area, spend an average of 7.1 hours seeking HIV-related care. HIV testing technology has also gradually improved in ways that make it easier to use. Originally HIV tests were carried out in batches in a laboratory, leading to delays of up to several days and requiring the patient to come back to receive the results. Thornton (2008) evaluates the older testing technology and randomizes the incentive to return and receive the results. So-called rapid HIV tests became available in developed countries in 2002 and several years later in sub-Saharan Africa. For these tests, a provider administers a saliva or blood test in the field and returns results within 20-30 minutes. Tests are 98 percent accurate and require a confirmatory test before reaching a definitive diagnosis. TLT surveyor implemented rapid HIV tests as part of this study. Self-administered home HIV tests are also increasingly available in sub-Saharan Africa. These tests are slightly less accurate but have important privacy advantages.

Evaluations of HIV testing interventions find that HIV testing has only a limited and contingent effect on risky sexual behavior (Thornton 2008, Delavande and Kohler 2012, Baird et al. 2014, Gong 2015). A meta-analysis of mostly non-experimental public health studies by Denison et al. (2008) is consistent with this finding. Boozer and Philipson (2000) and Gong (2015) note that testing should only influence individual risk optimization if test results lead people to update their HIV status beliefs. Most people have accurate priors about their HIV status, which limits the impact of testing on behavior. Secondly, the impact of new HIV status information on individual risk optimization is ambiguous and depends on the particular test result and the extent of updating about other epidemiological parameters like the transmission probability. Nonetheless, policymakers have sought to make HIV testing more universal. Thornton (2008) notes that policymakers have argued that HIV testing is the "missing weapon" in the fight against HIV/AIDS. The WHO and UNAIDS recently began recommending an "opt-out" approach to testing in countries with generalized epidemics (WHO 2007).

The Ministry of Health in Malawi provides free HIV testing and treatment through an existing network of primary health and antenatal clinics. These services are largely funded by international donors such as the Global Fund to Fight AIDS, TB, and Malaria. While nominally free, these services are inconvenient and require the recipient to spend several hours traveling and waiting. Pinto et al. (2013) surveyed recipients of HIV care in the Zomba District of Malawi, which is near the TLT study area. An average visit required 7.1 hours, including 3.3 hours traveling to and from the facility and 3.8 hours waiting in line. Patients also face uncertainty about whether providers will be able to see them on a particular day.

Stigma is another important non-monetary cost of seeking an HIV test. Merely seeking an HIV test may lead to rumors that a person is HIV positive or has been unfaithful (Sambisa et al. 2010, Berendes and Rimal 2011, Maughan-Brown and Nyblade 2014). Despite efforts to maintain confidentiality, observers in rural communities may notice whether someone has sought HIV-related care. This reputation cost magnifies the other costs of seeking a test. The implication that a test recipient is worried about his or her health is stronger in an environment where seeking a test is especially difficult. This dynamic may also explain why low-risk people do not test frequently in the status quo. Even if an HIV-negative test result provides a good signal, the act of seeking a test provides a bad signal in a regime with inconvenient opt-in testing. Policies such as provider-initiated testing may be particularly effective in terms of removing this aspect of stigma.

3.2 Marriage and HIV/AIDS

This subsection describes common marriage practices in the study area. Communities in the Balaka District of Southern Malawi predominately practice matrilineal kinship and matrilocal marriage, although compliance with these customs is not universal (Reniers 2008, Berge et al. 2014). Our sample is 81 percent Christian and 19 percent Muslim. 87 percent of people say religion is "very important" or "extremely important" in shaping daily life. The decision to marry rests with the couple rather than parents or other relatives (Kaler 2001, Kaler 2006). Polygamy is not common in our data. Among baseline-married men in our sample, less than 1 percent have multiple wives. As in other matrilineal settings, significant marriage payments are uncommon.[11]

Marriage in this setting is characterized by less formality and legalism than is present in a Western context. Marriage may involve a religious ceremony, a traditional ceremony, a legal registration, or simply the initiation of a cohabitation arrangement. In our data among baseline-married female respondents, 81 percent had a traditional ceremony, 16 percent had a religious ceremony, 8 percent registered their marriage legally, and 25 percent simply moved in together. These figures do not sum to one because people may have had multiple types of ceremonies. The lack of legal formality means that people may divorce more easily than in a Western context.

Women typically marry around age 20 and men marry around age 25 in our setting, so that the earnings potential of marriage-eligible men is largely observable, consistent with Bergstrom and Bagnoli (1993).

HIV prevalence is 6 percent in this setting. Although women are at a greater biological risk for contracting HIV, the age gap between men and women in the marriage market means that HIV prevalence is similar for men and women. This level of prevalence is high from a public health standpoint, but sufficiently low so that most people who are tested for HIV

[11] According to Meekers (1992), marriage payments (and other formalities) are less important in matrilineal societies because the marriage does not lead to a transfer of "ownership" of the wife's assets or the couple's children to the husband.

receive "negative" results.

Ethnographic studies in sub-Saharan Africa substantiate the idea that HIV risk has fostered adverse selection in the marriage market. Demographic research documents the correlation between HIV risk and the timing and match quality of marriage (Schatz 2005, Reniers 2008). Mukiza-Gapere and Ntozi (1995) and Watkins (2004) indicate that people believe that the primary threat of HIV/AIDS comes from the spouse. Mukiza-Gapere and Ntozi (1995) argue that in Uganda, "the fear of HIV/AIDS has 'poured cold water' on the institution of marriage", leading to marriage delays. According to Watkins (2004), women are concerned their husbands will "bring AIDS into the house", which will lead to suffering and make it difficult to care for children. People (particularly women) are conscious of this risk encourage each other to make marriage decisions strategically to avoid HIV.

The rise of HIV/AIDS in sub-Saharan Africa has also coincided with a trend toward later marriages in sub-Saharan Africa (Harwood-Lejeune 2001, Marston et al. 2009, Hosegood et al. 2009). Bongaarts (2007) also documents a cross-sectional correlation between HIV prevalence and marriage timing. Existing scholarship has treated marriage timing as exogenous and interpreted this relationship as evidence that marriage timing influences the spread of HIV (Magruder 2011). However our model provides an alternative interpretation for this pattern, in which the threat of HIV encourages people to delay marriage.

3.3 Data

This study uses data from the Tsogolo La Thansi (TLT) Panel Study, which was conducted in Balaka, Malawi from 2009 to 2011. The panel contains eight waves, each spaced four months apart. To enroll the sample, surveyors identified 1500 young women aged 15-25 from the area. These women were selected to be representative of this age group within the community. The study also includes a sample of 500 representative men of the same age group. Since women marry younger than men in this setting, the men in the data set are not marginal for marriage and therefore we do not consider them in this analysis. At each survey wave, participants were asked to name their three most recent sexual partners and provide a status update for each partner. Therefore we observe in each period whether the respondent is married. For both married and unmarried respondents, we observe the identity and some characteristics of the partner. In every wave, respondents were asked to take a pregnancy test. Compliance with this request was over 95 percent and most non-compliers were women who were obviously pregnant. The first survey wave occurred around June 2009. Waves were spaced approximately every four months for 2.5 years. The eighth and final survey wave occurred around October of 2011.

The data collection includes a partner sample. After each survey, participants were given tokens to provide to their partners and told that their partners could also enroll in the study. Partners who enrolled received a similar questionnaire to the index respondents. We do not rely extensively on the partner survey because selection into this sample is endogenous. Partners receive the same treatment assignment as the index respondents, which facilitates a screening mechanism through the survey. For treatment women, the partners who enroll in the study receive the same intervention, so that they are also subject to HIV testing. The willingness of a partner to participate in the survey may provide information about his risk profile.

The sample is evenly divided into three intervention arms. In the "treatment" arm, participants were offered HIV testing after every survey wave, for a total of eight tests. Surveyors provided rapid HIV tests, which yield results in several minutes. Surveyors had been trained to provide post-test counseling to respondents. The content of the counseling was contingent on the test result. Since the tests were offered after participants completed the survey, the information treatment contained in was measured at subsequent survey rounds. In the "control" arm, participants were only offered HIV tests after the Wave 8. The third "alternative treatment" arm was offered HIV tests only twice, immediately after Waves 4 and 8. As we describe below, this arm allows us to examine the impact of a one-off testing intervention.

Our analysis focuses on marriage and fertility. Marital status is measured in each wave, and is linked to a specific partner. For a respondent to be married, she must have a partner recorded as her husband in that period. Marriage is relatively informal in the study area, and in practice this variable equates to cohabitation. Around 44 percent of women are married at baseline. This figure rises to 63 percent by the end of the study. Divorce is possible in this setting, and could be an implication of the HIV testing intervention. Around 7 percent of respondents divorce over the study period. This rate is balanced across treatment and control groups (as we discuss further below), so that in practice we do not find treatment effects on divorce.

We measure pregnancy based on pregnancy test results that were collected in every wave. Respondents were asked to complete standard urine pregnancy tests. Compliance was over 90 percent, and most of the women who did not comply were obviously pregnant. We use pregnancy as a proxy for marital surplus in our tests of the model below. The premise for this interpretation is that children are an investment with a large up-front cost. Partners will only be willing to bear children if they perceive sufficient utility from the marriage. The link between surplus and fertility is obviously imperfect. One caveat is that pregnancies that precede marriage do not necessarily indicate high marital surplus. Such cases weaken but do not invalidate this proxy variable. As we show below, fertility patterns by HIV status and attractiveness are consistent with the surplus interpretation of pregnancy.

Surveyors assesses physical attractiveness on a five-point Likert scale with the following categories: 1 (Much more attractive than average), 2 (more attractive than average), 3 (Average), 4 (Less attractive than average), and 5 (much less attractive than average). Since only one respondent is classified in the final category, we combine this category with Group 4. Figure 1 shows the baseline frequency distribution of this variable. In the analysis below, we distinguish between people in Groups 1 and 2 (who are defined as attractive) and Groups 3, 4, and 5, who are defined as unattractive. One important caveat of this variable is that all surveyors were women. Therefore this variable captures female assessments of the attractiveness of respondents, rather than male assessments, which are arguably more relevant for the marriage market. This issue introduces measurement error, which should make it more difficult to find results.

The analog to physical attractiveness for men is income or earnings potential. While income is difficult to measure, completed education captures similar variation. Survey respondents indicate the level of completed schooling for each partner on a four-point scale, in which 0 indicates less than standard (primary) completion, 1 indicates standard completion, 2 indicates form (secondary) completion, and 3 indicates higher education. At baseline, around 1 percent of partners have did not complete standard, 40 percent completed primary, 52 percent completed form, and 7 percent completed higher education.

To examine the role of HIV status information, we estimate effects on perceptions of own and partner HIV status. The survey elicits the subjective probability that the respondent is currently HIV positive. To assess this and other subjective probability, surveyors used the "bean" methodology and took extra care to ensure that respondents understood the concept of probability and that they provided internally-valid responses. This variable ranges from 0 to 100 percent. 54 percent of people indicate zero probability that they are infected and 93 percent of people believe their infection probability is 50 percent or lower. In addition, the survey elicits the likelihood that the respondent's partner (not necessarily a spouse) has HIV on a five-point Likert scale. This variable is only available for respondents who have partners at the time of the interview, but the availability does not depend on whether the partner chooses to participate in the study. To simplify the interpretation of this variable, we create binary versions that equal 1 if the respondent believes her partner "may have HIV". Estimates using the full Likert scale are available from the authors.

3.4 HIV Testing Intervention

The HIV testing intervention was closely incorporated into the TLT Panel Study data collection. Researchers randomized respondents into three arms at baseline using a simple randomization. Participants in the treatment arm were offered HIV tests at the conclusion of every interview. Participants in the control arm were offered tests only after the final (eighth) wave. A third “alternative treatment” arm were offered tests twice, after Waves 4 and 8. Most of our analysis is based on a comparison of the treatment and control arms, however Section 4.4 below uses the alternative treatment arm to show that there are not impacts of a one-shot testing intervention on marriage and fertility.

Surveyors used “rapid tests” for HIV, which take around 30 minutes to complete. In order to administer these tests, surveyors underwent extensive certification in voluntary counseling and testing (VCT) from the Malawi Ministry of Health. Confirmatory tests were done for people who initially tested positive. The testing procedure involved both pre-test and post-test counseling related to HIV and HIV prevention. People who tested positive were told where to seek follow-up care. Respondents accepted the HIV tests around 80 percent of the time. Confidentiality is an important component of VCT, and so participants did not receive written documentation of their HIV status. As we describe above, proof of HIV status is not necessary for the existence of this testing regime to enable signaling and screening. Since the first tests were given at the end of Wave 1, we consider Wave 1 the baseline and Waves $2-8$ to be the follow-up.

All female participants were given tokens to share with up to three sexual partners, who could also enroll in the study. For participants in the treatment group, this component created a potential screening device, as a partner’s unwillingness to participate in the study may indirectly reveal that he is HIV-positive.

The interpretation of our estimates depends importantly on the outcome of the HIV test results. HIV testing enables signaling and screening, but only for people who test negative. Since prevalence is six percent, the sample of HIV-positive people is small, which limits the power to estimate effects in this group, despite the presence of theoretical predictions for HIV-positives. Isolating the HIV-negative people is useful since this step should strengthen our estimates if the model is correct. However HIV status is endogenous to treatment, and it is possible that the provision of HIV testing could have induced people to alter their risk exposure in a way that would influence HIV status. In practice, this phenomenon is not likely to be statistically important because the HIV sero-conversion rate is sufficiently low. In the treatment group, for whom we have multiple HIV tests over the sample period, only 11 out of 507 people change status over 2.5 years. It is unlikely that more than 1 or 2 of these people had behavior changes that affected their status because of the intervention, particularly since existing research suggests that HIV testing has only small effects on risky sexual behavior.

A second issue related to HIV testing is that data on HIV status are more complete for the treatment group than the control group. Since testing was voluntary, people can and did choose not to test on particular occasions. Since the control group was only offered tests once, there is a higher rate of coverage for the treatment group than the control group. For the treatment group, we observe the HIV status at least once for 97 percent of participants. For the control group we observe the status for 75 percent of participants. Statistically, it is likely that nearly all of the non-testing people in the control group are HIV negative. Among testers, the prevalence is 10.5 percent in the treatment group and 4.9 percent in the control group. To match the prevalence in the treatment group, prevalence among non-testers would need to be around 20 percent. Therefore classifying non-tester as HIV negative for the purpose of the analysis reduces measurement error compared to the alternative of classifying them as HIV-positive. Estimates are robust to varying this classification of non-testers.

4 The Impact of Intensive HIV Testing

This section estimates the impacts of the HIV testing intervention. To support the adverse selection interpretation, we first show that the intervention provides information about partner HIV status but (for the most part) does not update beliefs about own HIV status. Next we show effects on marriage and fertility both overall, and by baseline marital status, HIV status, and attractiveness.

4.1 Identification and Estimation

Our primary specification pools the follow-up waves (Waves $2-8$) to estimate a combined average treatment effect. We estimate the following equation:

$$Y_{it} = \beta_0 + \beta_1 T_i + \beta_2 Y_i^b + X_i^{b\prime} \beta_3 + \delta_t + \varepsilon_{it} \tag{12}$$

In this equation, Y is the outcome of interest, marriage or fertility, T is an indicator for the assignment to the treatment group, Y^b is the dependent variable at baseline, and X^b is a vector of other baseline characteristics, which we include to increase the precision of the estimates. δ_t is a set of wave indicators. We estimate Equation (12) by OLS and cluster the standard errors by subject.

Being offered an HIV test in every wave is the "treatment" for this analysis. Therefore assignment to and receipt of treatment are equivalent. β_1 identifies the average treatment effect on the treated (ATT) under the stable unit treatment value (STUVA) and ignorability assumptions. General equilibrium effects on the marriage market are a potential threat to the STUVA. However the treatment group represents only around 10 percent of local marriage market participants, minimizing this concern.

One econometric issue for our analysis is that the treatment and control arms are not balanced by age. While all respondents range from 15 to 25 at baseline, control respondents are an average of 0.6 years younger than treatment respondents. This imbalance is apparently due to chance, since other orthogonal characteristics are balanced. Since marriage and fertility increase with age in this range, the age imbalance could spuriously indicate higher marriage and pregnancy propensities for the treatment group. We address this issue by employing entropy weights throughout the analysis. Entropy weights, which are similar to inverse propensity weights, balance the data so that the treatment and control arms have the same mean, variance, and skewness (Hainmueller 2012, Hainmueller and Xu 2013).

Table 1 provides summary statistics for the treatment and control arms conditional on age reweighting.[12] Demographic and socioeconomic variables appear balanced, including tribe, religion, HIV prevalence, school enrollment, and employment. The household asset index is a standardized sum of indicators that the household has a durable roof, a durable floor, electricity, a television, a telephone, and an improved toilet. The table also shows that treatment and control respondents have similar levels of future orientation and subjective mortality risks. Balance for these parameters is important because the internal discount rate, δ, determines marriage timing in our model. Finally, the table shows that marriage and pregnancy rates do not differ significantly across treatment arms at baseline. By including the baseline dependent variable, Equation (12) also controls for any remaining imbalance in these outcomes.

After showing average treatment effects, we estimate heterogenous effects by HIV status and baseline attractiveness for baseline-unmarried respondents. Columns $1-3$ of Table 2 distinguish between HIV-positive and HIV-negative respondents. We do not oversample HIV-positive respondents, and the unmarried subsample only includes 37 HIV-positive women. As anticipated, the HIV-positive and HIV-negative groups differ in several important ways. HIV-positive women are older and have lower socioeconomic status. They also perceive significantly higher mortality risk. Attractiveness does not differ significantly by HIV status, which is consistent with the theoretical assumption that health and attractiveness are independent.

Columns $4-6$ of Table 2 further limit the sample to HIV-negative baseline-unmarried respondents (consistent with subsequent regressions) and cut by attractiveness. Attractiveness is not correlated with demographic characteristics, however attractive women have higher socioeconomic status in several dimensions. Attractiveness may be correlated with SES if wealthy respondents spend more on clothing and grooming. Attractive women also "think about the future" to a greater degree than unattractive women. These correlations may confound our estimates if attractiveness interactions spuriously capture interactions with SES

[12]Unweighted versions of all results are available from the authors. The unweighted version of Table 1 shows that marriage and fertility are significantly higher in the treatment group.

or time preferences. We discuss these issues further below.

4.2 The Impact on Information About Partner Quality

The model describes the effects of making health observable. To do that, our treatment, offering high-frequency free HIV testing, must increase the frequency of HIV testing and change the market's perception of the health status of treated people. We present several pieces of evidence consistent with these predictions.

First, Figure 8 shows that the treatment more than doubles the likelihood that women and their partners were tested in the previous 4 months, increasing it from 30 to 70 percent for women and from 25 to 48 percent for their partners. This is consistent with the treatment reducing the cost of HIV testing.

Second, Table 3, Panel A shows that baseline beliefs are positively correlated with HIV status, as the baseline belief of being HIV positive is 43 percentage points lower for HIV negative women than for HIV positive ones. Out of all HIV positive women in our data, about half are certain to be HIV positive and about 80 percent think they are at least 50 percent likely to be HIV positive. Moreover, this Panel shows that being offered high-frequency HIV testing does not change HIV-negative treated women's belief of their HIV status, despite the fact that almost all women are tested at least once in the 2.5 years of the intervention. This suggests that women do not get tested to find out their HIV status, as, if it were the case, we would find them to update their belief after being tested.[13] That is, most treated women choose to be tested multiple times despite having correct beliefs of being HIV negative. This is consistent with HIV testing being used to *signal* rather than to find out one's HIV status.

One challenge for the estimates in Panel A is that people may update their beliefs at different times within the two-year follow-up period. Estimates average across waves in which

[13] The increase in belief of being HIV positive after testing positive is driven by X (few) women, who must have been unaware of being infected with HIV. Nevertheless, the remaining HIV+ - X women, that is, the majority of HIV positive women, do not change their belief after finding out their HIV status.

new information has and has not been revealed. Panel B of Table 3 illustrates the impact on perceptions of partner health using an alternative source of variation. We implement individual fixed effects regressions within the treatment group and estimate the effect of receiving an HIV-negative own or partner test result in the previous wave, compared to the counterfactual of receiving either a positive result or no test result. These results lack a causal interpretation because time-varying unobservables could influence both perceptions and the decision to test. The fixed effects control for time-constant unobservable factors. However this approach isolates more effectively the instances in which people receive *new* information from the intervention.

In Column 1, receiving a negative test result reduces the subjective perception being HIV positive by 6.6 percentage points. However Column 2 shows that this effect is weaker and insignificant if we limit the sample to people with partners and control for partner status information. These findings are consistent with Panel A, which shows only a limited effect of HIV status information on perceptions of own status. In Column 3 of Panel B, receiving a negative test result from a partner leads to a significant reduction in the belief that the partner may be HIV positive.

Figure 9 illustrates this pattern graphically by showing the perceived partner HIV status for women whose partners ever test positive, do not test, and test negative at least once. Women whose partner tests negative at least once are less likely to believe their partner is HIV positive than women's beliefs that their partner is HIV positive correlate with the partner's testing behavior and outcomes: 83 percent of women whose partners test positive believe the partners may have HIV, compared to 49 percent for women whose partners do not test, and 25 percent for women whose partners test negative at least once.

We cannot tell to what extent these patterns are caused by selection (e.g., HIV positive are less likely to accept to get tested) and to what extent by belief updating (e.g., if partners tests negative, women become more confident their partner is negative and faithful). Nevertheless, both selection and updating are consistent with testing as a signaling device.

Given that people are likely to have better than average information about their partner's type, we interpret changes in these beliefs as a lower bound of the effect of the treatment on the marriage market's beliefs regarding the health type of its participants.

Third, Figure 12 provides some additional evidence consistent with a signaling use of HIV testing: among HIV-negative women, women with the highest confidence of being HIV negative at baseline test more frequently than women who are less certain of being HIV negative. If testing were used mainly to find out one's status, then we would find that women who are less certain of being HIV negative (perhaps because they behave in ways that increase their likelihood of becoming infected) would test more frequently.

Lastly, we would like to observe how the treatment changes the market's beliefs regarding people's health. Unfortunately, we do not observe the belief of marriage market participants about other people's HIV status, unlike for women's beauty. However, we have information on women's beliefs about their partner's HIV status.

Taken together, the evidence presented in this section is consistent with the hypothesis that our treatment lowers the cost of HIV testing and that testing has a signaling role in the marriage market.

4.3 Treatment Effects on Marriage and Fertility

Before testing the model's prediction, Table 4 shows the estimates of the ATT effects on marriage and pregnancy for the entire sample (Panel A) and by both HIV status (panel B) and baseline marital status (Panel C). The likelihoods of being married and pregnant have statistically significant increases of 0.035 and 0.021, that is, 6 and 12 percent increases. Based on this evidence, we can conclude that the treatment increases average surplus, and therefore, utility, for our subjects. Consistent with our model, the estimates are larger and more precisely estimated for HIV negative women and for unmarried women. Nevertheless, note that the estimated ATT effects do not vary statistically by health or marital status.

Table 5 presents our main results by considering women unmarried at baseline only, since the model predicted changes in marriage and fertility for this group only. Panel A shows statistically significant increases in the likelihoods of being currently married and currently pregnant for HIV negative women and statistically insignificant decreases in these outcomes for HIV positive women. Both sets of signs are consistent with the model predictions that the treatment should induce healthy women to marry earlier and have a higher marital surplus, while it should not change marriage rates and surplus for unhealthy women. As before, the two sets of estimates do not statistically differ by women's HIV status. However, since we have only 73 HIV positive women in our sample, it is not surprising that their ATT effects are imprecisely estimated.

Panel B considers HIV negative unmarried women only, and estimates ATT by attractiveness. As predicted by the model, the ATT effect on marriage is statistically larger for attractive than unattractive women, while the ATT effects on fertility are positive for both groups and not statistically different from each other.

Since health and attractiveness are correlated with various determinants of marriage and fertility, Columns 2 and 4 of Table 5 estimate a version of Equation (12) that further interacts the treatment dummy by three standardized indices of demographic, socio-economic, and time preference variables, as in Katz et al. (2001). Comparing across columns shows that the estimated ATT are stable regardless of which specification we use and that, therefore, the positive ATT effects are not caused by attractive and healthy women having different demographic and socio-economic characteristics or being more forward-looking.

Figure 10 shows this pattern graphically by plotting the marriage rate for baseline-unmarried women in the treatment and control groups by attractiveness. In the control group, marriage is substantially delayed for attractive women compared to unattractive women. However in the treatment group, this delay disappears and both types marry a similar pace. This pattern is reassuring because it shows that the attractiveness interaction is driven by a changes in the treatment group relative to a theoretically-anticipated pattern in the control group.

An additional implication of our model is that the HIV testing intervention should not lead to changes in matching on observable attractiveness. Estimating this effect is challenging because we only observe husband characteristics for women who actually marry. Nonetheless, Figure 11 shows that among married women at Wave 8, treatment and control participants have husbands with nearly the same levels of education. This pattern suggests that our assumptions about assortative matching within HIV categories are reasonable, since these assumptions lead us to this implication.

4.4 The Effect of a One-Shot HIV Testing Intervention

Few studies have examined the impact of HIV testing on marriage. One recent example comes from Beegle et al. (2015), who find no effect of a one-off testing intervention on marriage and fertility in Malawi. Their study, like the others in this literature that find no effects on risky sexual behavior, typically evaluate the effect of single HIV testing interventions. We argue that, in order for not being tested to send a negative signal about one's health, HIV testing must be offered (and turned out) multiple times: refusing to be tested multiple times sends a stronger signal than not being tested only once.

The TLT design allows us to examine the impact of a single test through an alternative treatment arm that was offered HIV tests twice, in Waves 4 and 8. To estimate the effect of being offered HIV testing only once, we compare marriage and fertility for this alternative treatment arm and the control group including only Waves 4 to 8. In this approach, Wave 4 is the baseline and Waves $5-8$ are the follow-up.[14] This approach is not entirely analogous to our earlier experiment because there are fewer follow-up waves.[15]

Table 6 reproduces the estimates from Table 5 for this alternative treatment. Panel A cuts by HIV status and Panel B cuts by attractiveness conditional for the HIV-negative

[14] Recall that the test is offered after the survey is completed. Therefore, a test given in Wave 4 cannot affect behavior and responses in Wave 4, but it can for Wave 5.

[15] It is not feasible to compare the original (intensive) testing intervention to the one-shot testing intervention because this alternative approach requires us to start at Wave 4, by which time, the original treatment arm had already received three HIV tests.

sample. Both panels show no significant effects of the one-shot testing intervention on marriage and fertility. The patterns that we previously observed, in which HIV-negative and attractive women responded differentially, are no longer apparent. Main effect estimates that are analogous to Table 4 are available from the authors and also show no impact. These results suggest that the frequency of testing in our main intervention is a key reason why HIV testing facilitates marriage.

5 Conclusion

Adverse selection is an important phenomenon in various product and labor markets that feature asymmetric information. Little or no research to date has examined the implications of this model in the marriage market. We combine a simple assortative matching model with the randomized evaluation of a novel intervention to reduce asymmetric information related to HIV risk among partners. Women aged 15 to 25 in the treatment arm, as well as their partners, were offered free opt-out HIV tests every four months for 2.5 years. We show that this intervention increased the probability of marriage and pregnancy by 3.5 percentage points and 2.1 percentage points, respectively, for these women. Estimates are stronger for HIV-negative and baseline-unmarried women. We also find that women who are physically attractive contribute the most to this effect. Under assortative matching, these women have access to the most upwardly mobile men, so that the stakes for selecting an HIV-positive partner are greater.

Researchers have found small and contingent effects of HIV testing on risky sexual behavior, which has limited the potential role of HIV testing in the eyes of some policymakers. A key distinction between this evaluation and others is the frequency with which tests were offered. Participants and their partners received up to eight HIV tests at regular intervals for 2.5 years in this intervention, altering perceptions of the cost of testing and enabling people to signal and screen. These benefits of testing make less sense for interventions that offer a single test. To support this interpretation, we show that an alternative one-shot testing intervention in the same population had no significant effects on marriage and fertility, even among the subgroups who respond the most to repeated testing. Despite the lack of a response for risky sexual behavior (we also find insignificant effects on these outcomes), an impact on marriage is relevant for the epidemiology of HIV because discordant spouses are a major source of new HIV infections.

While unobservable partner quality does not have a strong influence on education choices in this setting, it may matter elsewhere. Women and girls have a marriage market incentive to marry early in settings where men prize (but cannot observe) virginity. In addition, our findings suggest that the HIV/AIDS epidemic has contributed to delays in marriage in endemic settings. Without the epidemic, people might marry sooner, which could discourage schooling for younger girls.

Table 1: Baseline Characteristics by Treatment Status

	Full Sample		
	Treatment	Control	P-value
	(1)	(2)	(3)
Demographics			
Age	19.8	19.8	1.00
Attractiveness	3.54	3.59	0.21
Ngoni Tribe	0.38	0.38	0.99
Yao Tribe	0.25	0.26	0.83
Lomwe Tribe	0.19	0.16	0.15
Catholic	0.33	0.32	0.71
Protestant	0.49	0.49	0.89
Muslim	0.18	0.19	0.78
HIV positive (endline)	0.10	0.08	0.14
Socioeconomic Status			
Enrolled in school	0.36	0.40	0.14
Employed full-time	0.18	0.20	0.43
Any savings	0.17	0.13	0.12
Household asset index	-0.02	0.06	0.16
Preferences and Perceptions			
Thinks about future	3.12	3.19	0.28
Subjective 5-year mort. risk	0.34	0.33	0.74
Subjective probability HIV positive	0.12	0.10	0.17
Worried about HIV	1.04	1.03	0.85
Outcomes			
Married	0.43	0.46	0.26
Pregnant	0.15	0.12	0.15
Tested within 4 months	0.22	0.17	0.07*
Partner may be HIV positive	0.48	0.44	0.44
Observations	500	507	-

Note: All means are weighted for age balance. To compute p-values, we regress each variable on treatment in Wave 1 and cluster standard errors by respondent. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Baseline Characteristics by HIV Status and Attractiveness

	Unmarried Sample			Unmarried / HIV Negative Sample		
	HIV−	HIV+	P-value	Attractive	Unattractive	P-value
	(1)	(2)	(3)	(4)	(5)	(6)
Demographics						
Age	18.4	21.2	0.00***	18.3	18.4	0.75
Ngoni Tribe	0.35	0.40	0.56	0.36	0.35	0.98
Yao Tribe	0.25	0.21	0.60	0.23	0.27	0.36
Lomwe Tribe	0.18	0.14	0.47	0.19	0.19	0.98
Catholic	0.38	0.56	0.03**	0.36	0.40	0.31
Protestant	0.48	0.27	0.01***	0.50	0.43	0.17
Muslim	0.16	0.17	0.84	0.15	0.17	0.59
Attractiveness	3.61	3.38	0.15	4.14	2.96	0.00***
Socioeconomic Status						
Enrolled in school	0.68	0.23	0.00***	0.76	0.59	0.00***
Employed full-time	0.07	0.36	0.00***	0.04	0.10	0.02**
Any savings	0.12	0.22	0.17	0.10	0.14	0.26
Household asset index	0.30	-0.14	0.01**	0.58	-0.05	0.00***
Preferences and Perceptions						
Thinks about future	3.32	3.04	0.15	3.54	3.04	0.00***
Subjective 5-year mort. risk	0.31	0.43	0.04**	0.33	0.30	0.24
Subjective probability HIV+	0.08	0.34	0.00***	0.08	0.08	0.65
Worried about HIV	0.73	1.18	0.03**	0.70	0.73	0.78
Outcomes						
Pregnant	0.05	0.07	0.55	0.05	0.04	0.68
Tested within 4 months	0.15	0.05	0.01**	0.14	0.15	0.73
Partner may be HIV positive	0.40	0.75	0.00***	0.37	0.45	0.23
Respondents	552	37	-	303	249	-

Note: To compute p-values, we regress each variable on HIV status or attractiveness in Wave 1 and cluster standard errors by respondent. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: The Impact of HIV Testing on HIV Perceptions

	Respondent pr(HIV+)		Partner May Have HIV
	(1)	(2)	(3)
A: Estimates by HIV Status			
Treatment · HIV Negative	0.0048	0.011	-0.049**
	(0.0091)	(0.010)	(0.020)
treatment · HIV Positive	0.19***	0.19***	0.055
	(0.060)	(0.066)	(0.066)
Control mean (HIV negative)	0.14	0.15	0.34
Control mean (HIV positive)	0.57	0.59	0.71
Observations	6051	4281	4281
Sample	S1	S2	S2
B: Individual Fixed-Effects Estimates			
Respondent tests negative	-0.066***	-0.023	-0.035
	(0.018)	(0.021)	(0.038)
Parter tests negative	-	-0.0097	-0.11***
		(0.015)	(0.035)
Control mean	0.19	0.21	0.36
Observations	3508	2087	2087
Sample	S3	S4	S4

Note: Clustered standard errors appear in parentheses. The S1 sample (Column 1 of Panel A, our main sample for marriage and fertility estimates below) covers Waves $2-8$ and includes all respondents. The S2 sample (Columns $2-3$ of Panel A) covers Waves $2-8$ for respondents with partners. The S3 sample (Column 1 of Panel B) covers Waves 1-8 for the treatment group only. The S4 sample (Columns $2-3$ of Panel B) covers Waves $1-8$ for respondents with partners in the treatment group. All estimates control for the wave fixed effects. Panel A estimates control for the baseline dependent variable. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: The Impact of HIV Testing on Marriage and Fertility

	Currently Married	Currently Pregnant
	(1)	(2)
A: Main Effects		
Treatment	0.035**	0.021**
	(0.017)	(0.010)
Control mean	0.55	0.13
B: Estimates by HIV Status		
Treatment · HIV Negative	0.037**	0.022**
	(0.018)	(0.011)
Treatment · HIV Positive	-0.013	0.011
	(0.073)	(0.036)
Control mean (HIV negative)	0.54	0.14
Control mean (HIV positive)	0.61	0.14
Equality of coeffs. (p-value)	0.51	0.78
C: Estimates by Baseline Marital Status		
Treatment · Unmarried	0.056**	0.031**
	(0.028)	(0.013)
Treatment · Married	0.010	0.016
	(0.018)	(0.015)
Control mean (Unmarried)	0.21	0.12
Control mean (Married)	0.93	0.17
Equality of coeffs. (p-value)	0.17	0.47
Observations	6048	6048

Note: Clustered standard errors appear in parentheses. All regressions include Waves $2-8$ and reweight to balance by age. Estimates control for the baseline dependent variable and wave indicators. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Estimates for Women Who are Unmarried at Baseline

	Currently Married		Currently Pregnant	
	(1)	(2)	(3)	(4)
A: Estimates by HIV Status				
Treatment · HIV Negative	0.052*	0.052	0.034**	0.033**
	(0.030)	(0.035)	(0.014)	(0.015)
Treatment · HIV Positive	-0.077	-0.040	-0.016	-0.0076
	(0.13)	(0.13)	(0.056)	(0.060)
Equality of coefficients (p-value)	0.33	0.51	0.39	0.52
Observations	3398	3398	3398	3398
B: HIV-Negative Estimates by Attractiveness				
Treatment · Attractive	0.10***	0.12***	0.044**	0.049**
	(0.038)	(0.044)	(0.018)	(0.020)
Treatment · Not Attractive	-0.015	-0.0084	0.013	0.015
	(0.048)	(0.047)	(0.021)	(0.021)
Equality of coefficients (p-value)	0.06	0.03	0.27	0.21
Observations	3174	3174	3174	3174
Control for:				
Treatment · Demographics	-	Yes	-	Yes
Treatment · SES	-	Yes	-	Yes
Treatment · Time prefs.	-	Yes	-	Yes

Note: Clustered standard errors appear in parentheses. All regressions include Waves $2-8$ and reweight to balance by age. Estimates control for the baseline dependent variable and wave indicators. Demographic controls include tribe, religion, and age. SES controls include indicators for school enrollment, employment, durable roof, durable floor, electricity, telephone ownership, and television ownership. Time preference controls include future orientation and subjective mortality risk within 1, 5, and 10 years. For Columns 2 and 4, we interact Treatment with the first principal component of the baseline variables within each group. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: The Impact of an Alternative Single-Test Intervention on Unmarried Women

	Currently Married		Currently Pregnant	
	(1)	(2)	(3)	(4)
A: Estimates by HIV Status				
Treatment · HIV Negative	-0.024	-0.018	-0.0085	-0.0074
	(0.031)	(0.030)	(0.019)	(0.019)
Treatment · HIV Positive	-0.019	-0.019	0.034	0.036
	(0.065)	(0.067)	(0.049)	(0.048)
Equality of coefficients (p-value)	0.94	0.99	0.39	0.38
Observations	1617	1617	1617	1617
B: HIV-Negative Estimates by Attractiveness				
Treatment · Attractive	-0.027	-0.012	-0.026	-0.020
	(0.035)	(0.035)	(0.025)	(0.025)
Treatment · Not Attractive	0.0093	0.014	0.013	0.011
	(0.049)	(0.048)	(0.030)	(0.030)
Equality of coefficients (p-value)	0.54	0.66	0.32	0.45
Observations	1479	1479	1479	1479
Control for:				
Treatment · Demographics	-	Yes	-	Yes
Treatment · SES	-	Yes	-	Yes
Treatment · Time prefs.	-	Yes	-	Yes

Note: Clustered standard errors appear in parentheses. All regressions include Waves $5-8$ and reweight to balance by age. Estimates control for the Wave-4 dependent variable and wave indicators. Demographic controls include tribe, religion, and age. SES controls include indicators for school enrollment, employment, durable roof, durable floor, electricity, telephone ownership, and television ownership. Time preference controls include future orientation and subjective mortality risk within 1, 5, and 10 years. For Columns 2 and 4, we interact Treatment with the first principal component of the Wave-4 variables within each group. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

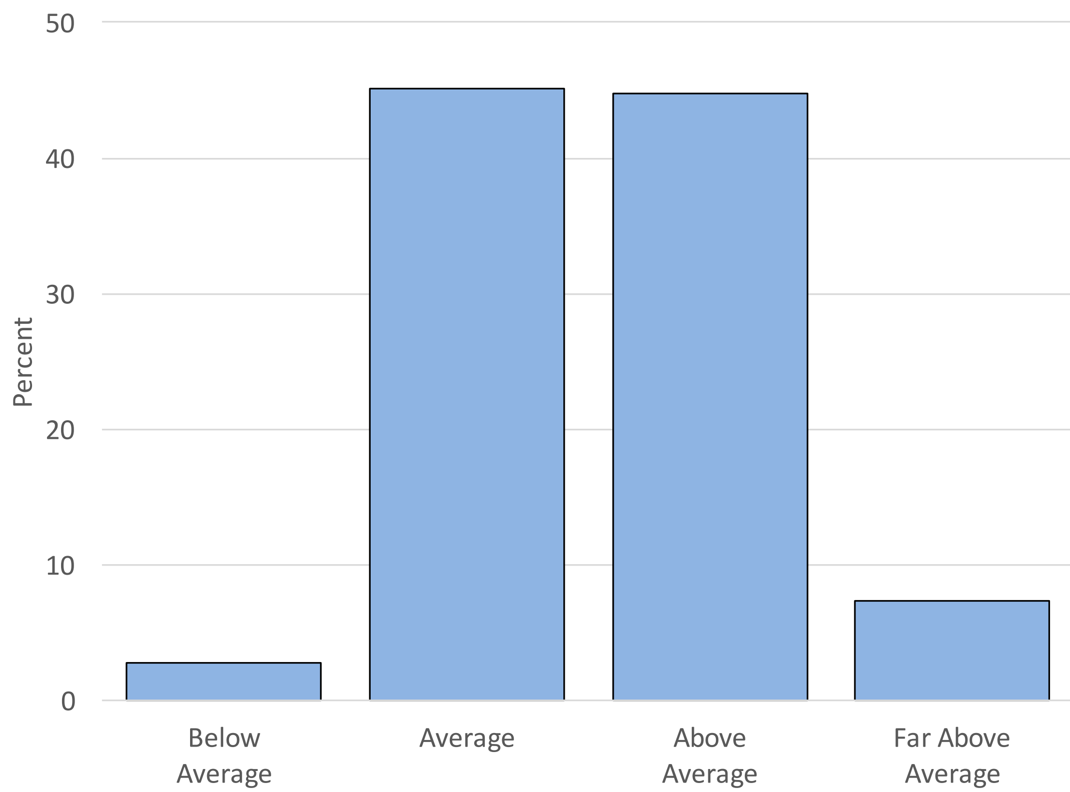

Figure 1: The Distribution of Baseline Attractiveness

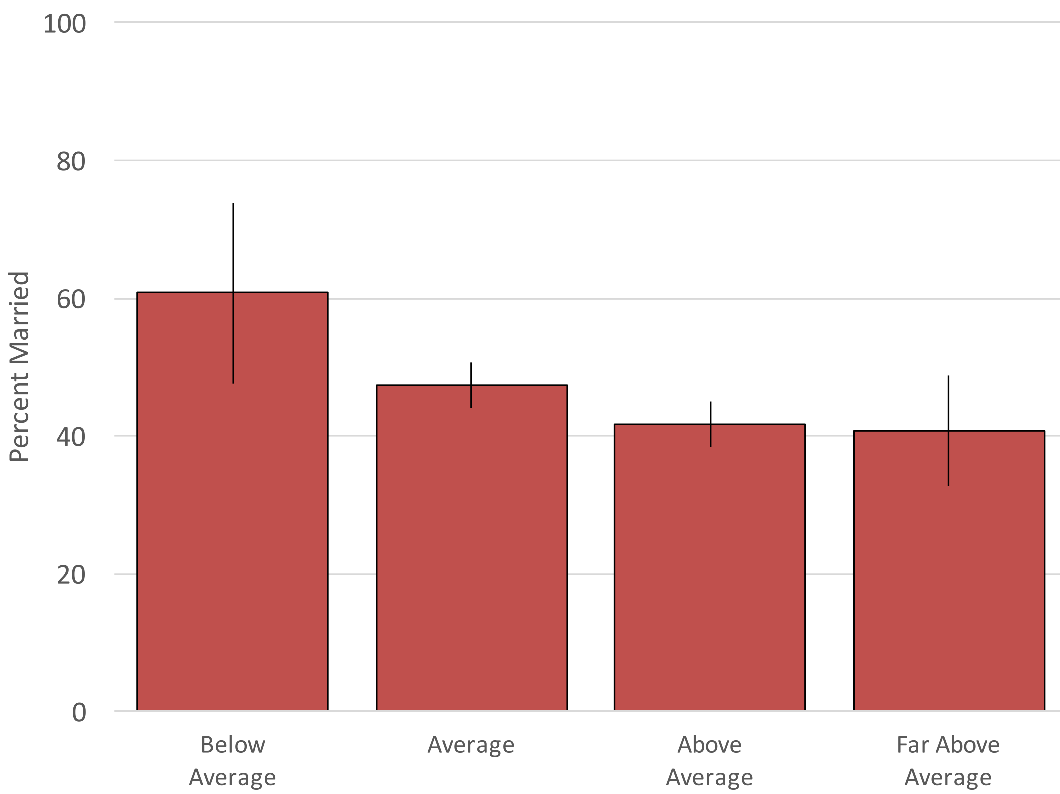

Figure 2: The Baseline Marriage Rate by Attractiveness

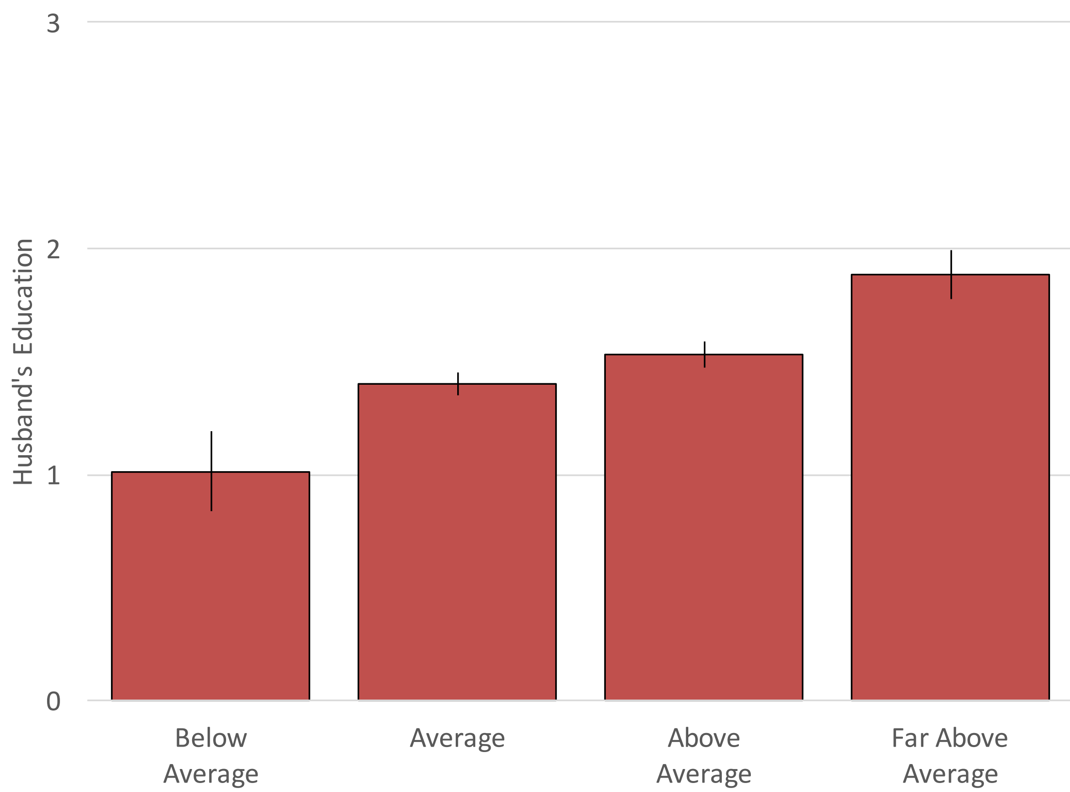

Figure 3: Baseline Husband's Education by Attractiveness

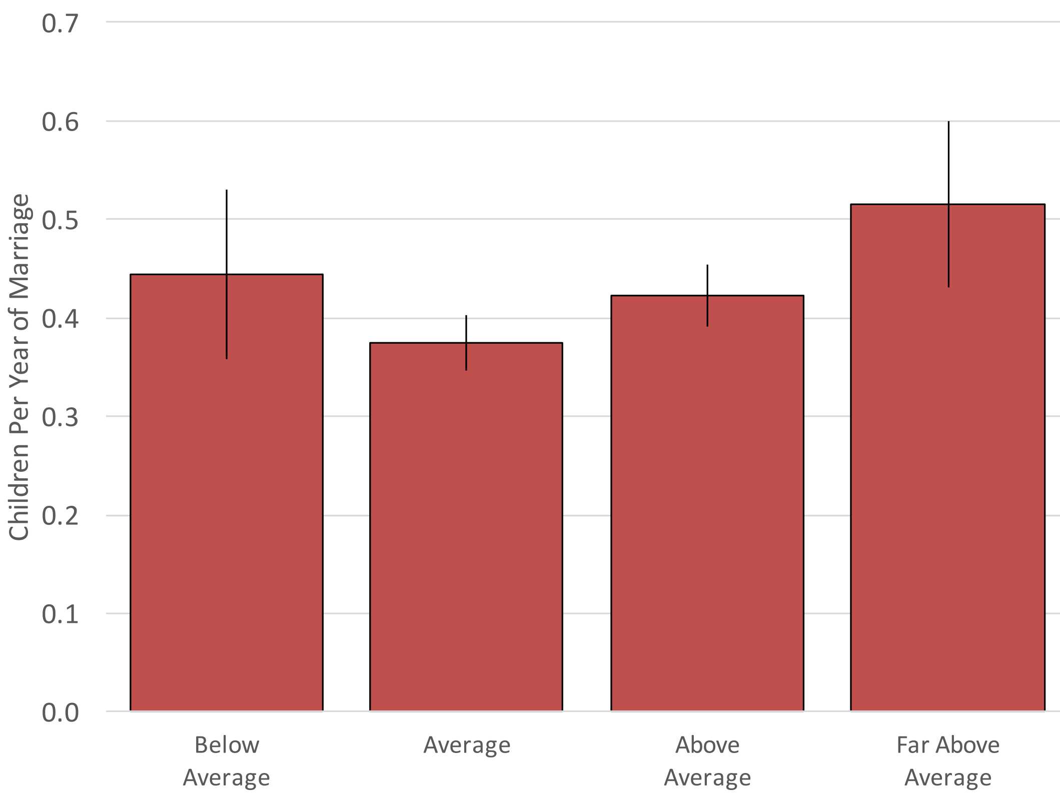

Figure 4: Children per Year of Marriage by Attractiveness

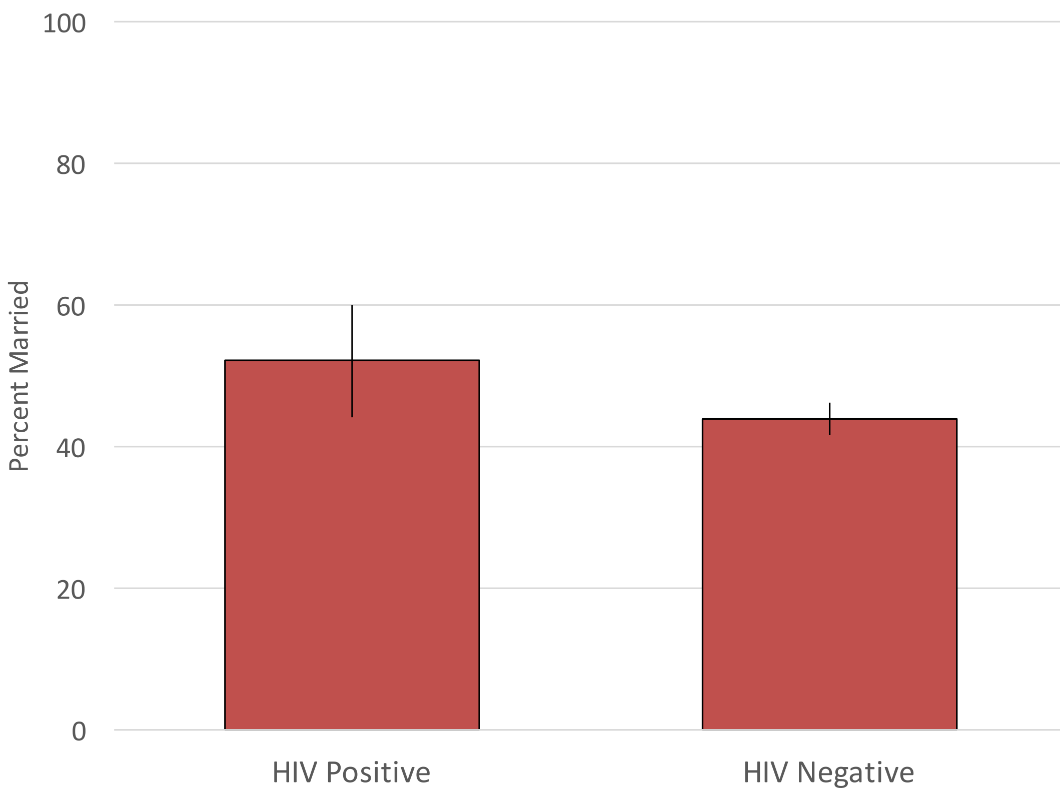

Figure 5: The Baseline Marriage Rate by HIV Status

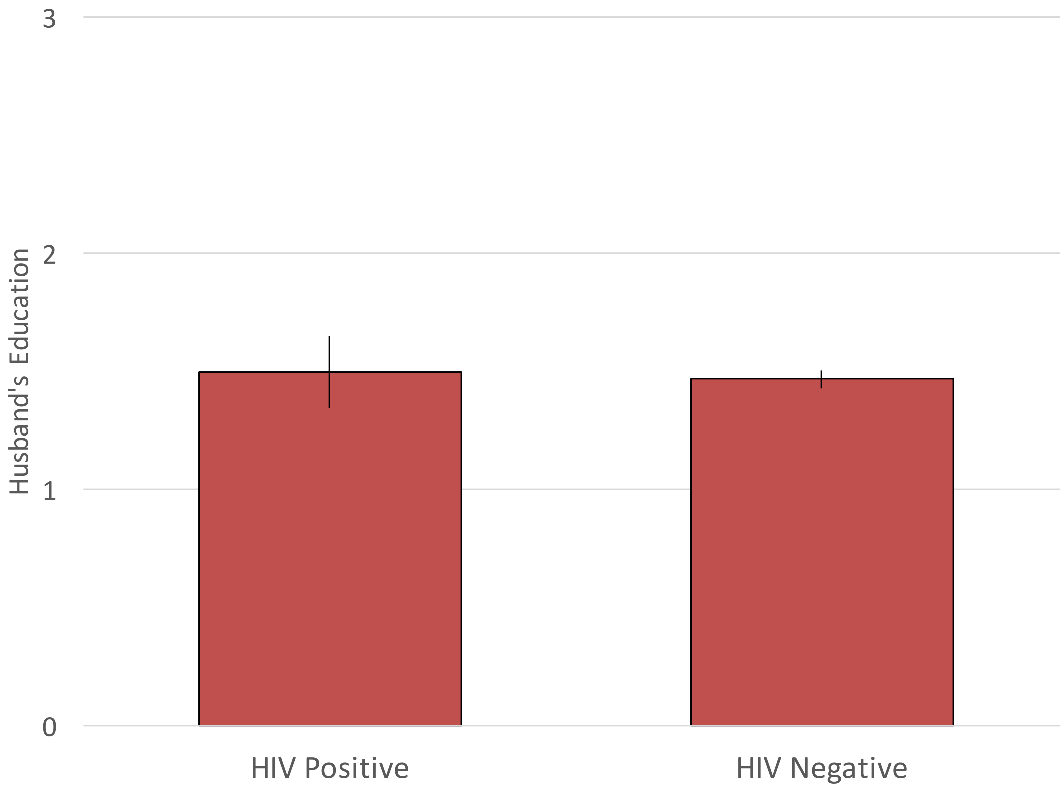

Figure 6: Baseline Husband's Education by HIV Status

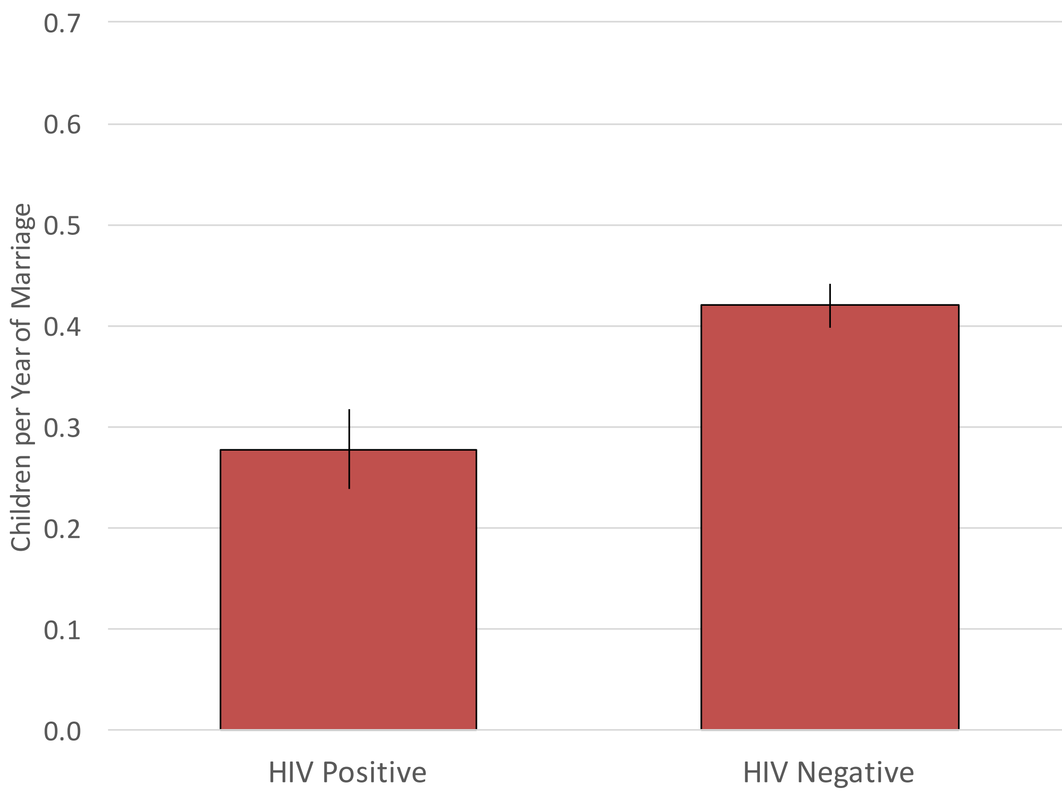

Figure 7: Baseline Children per Year of Marriage by HIV Status

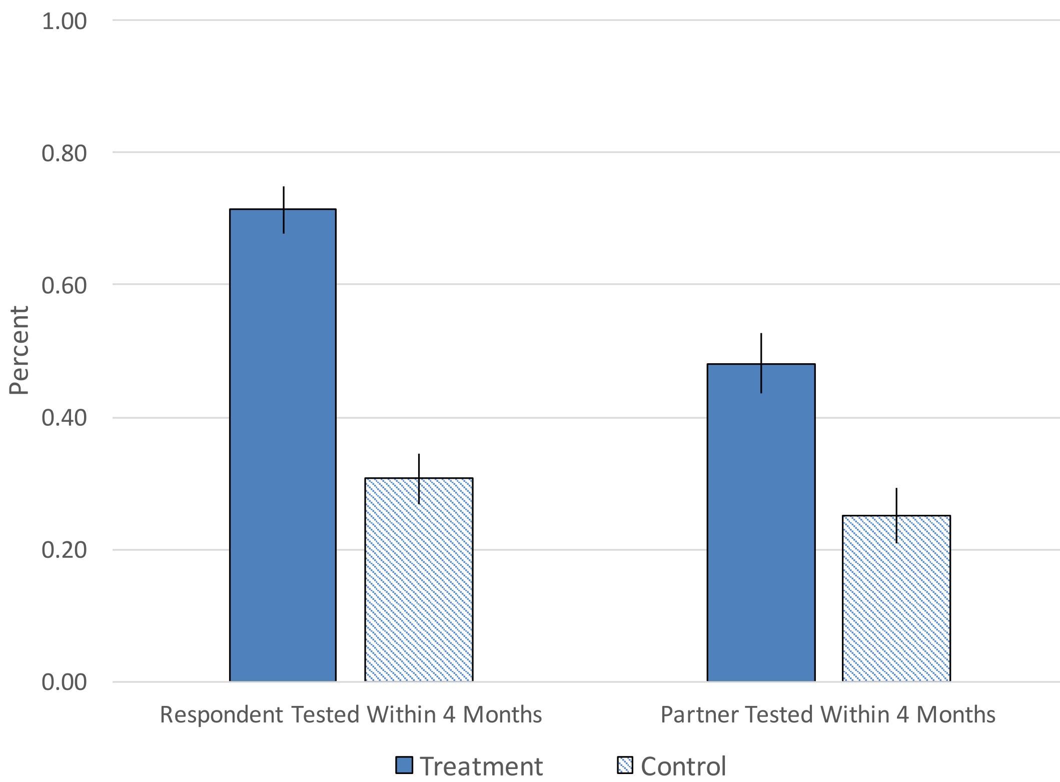

Figure 8: Probability of Testing within Four Months by Treatment Arm

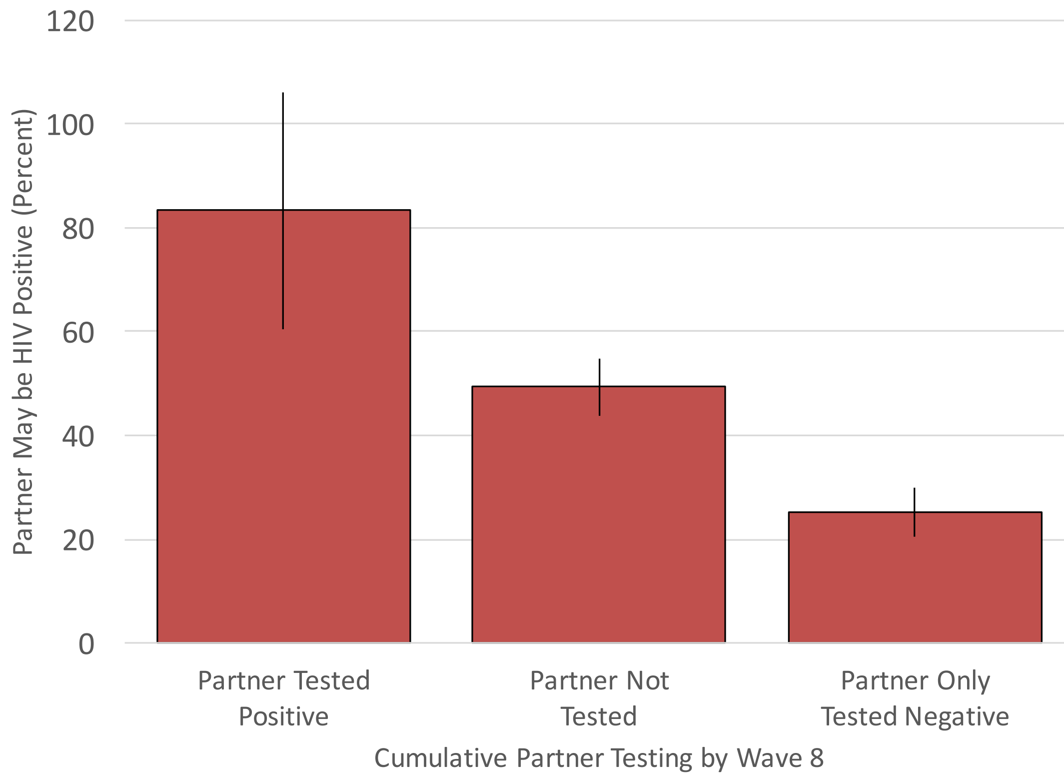

Figure 9: Perceptions of Partner HIV Status at Wave 8

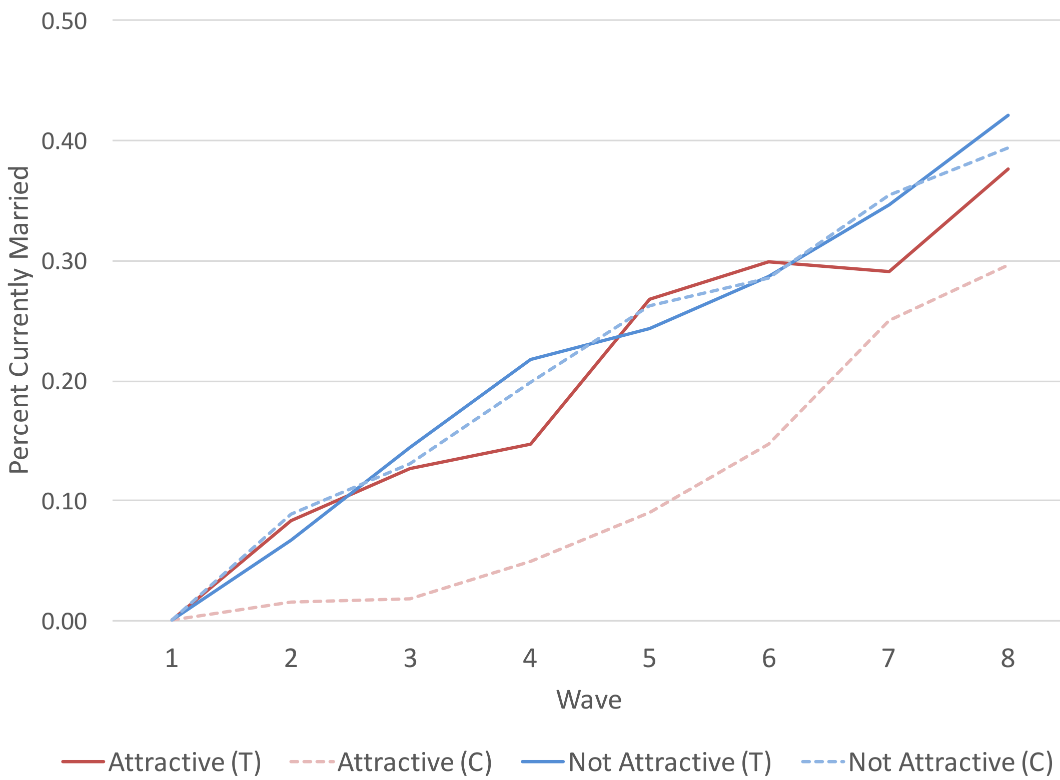

Figure 10: The Marriage Rate by Attractiveness for HIV-Negative Baseline Singles

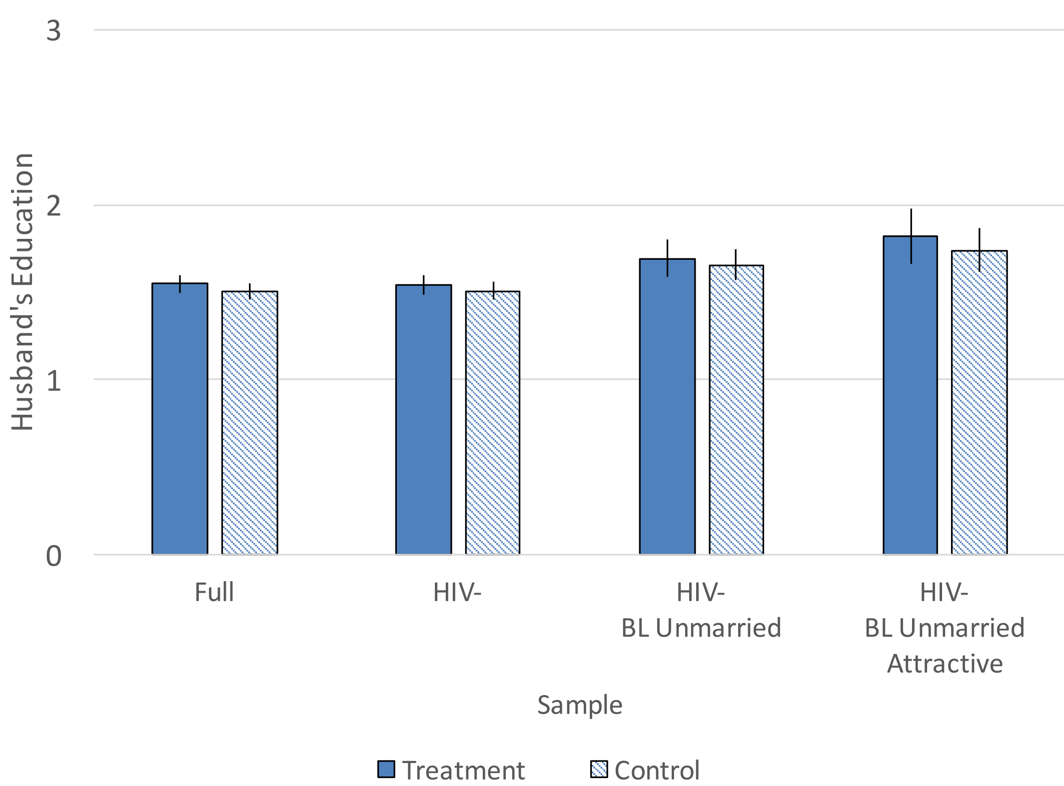

Figure 11: Husband's Education for Married Women by Treatment Status at Wave 8

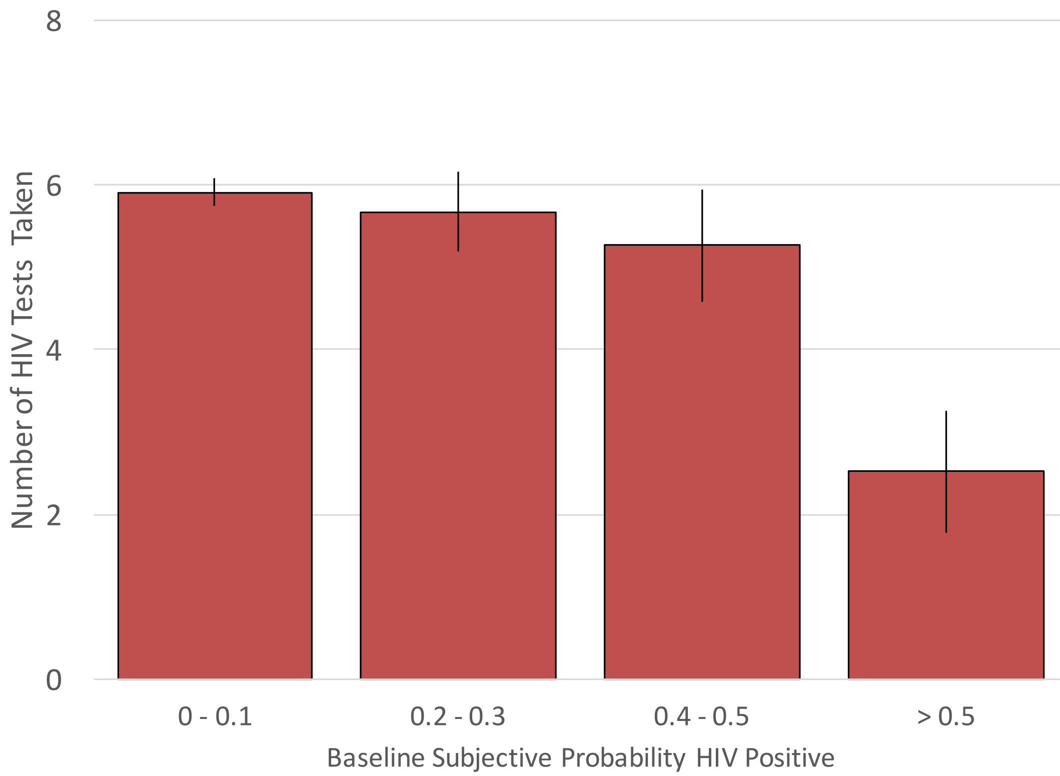

Figure 12: The Baseline Subjective Probability HIV+ and Subsequent Testing Frequency

6 Appendix: Alternative Assumptions

This section considers the implications of alternative assumptions about the relative values of attractiveness and health, divorce, and the correlation between attractiveness and health.

6.1 Relative Value of Traits

For the following relationship to hold, $S^{ab}_{hh} > S^{ab}_{hl} = S^{ab}_{lh} > S^{ab}_{ll} > 0$, each trait must have a similar effect on marriage surplus. However, if health were a larger determinant of surplus than attractiveness, marrying an unattractive but healthy person, a type lh, would yield a higher surplus than marrying an attractive but unhealthy person, a type hl and, therefore, the surpluses would be ranked as follows: $S^{ab}_{hh} > S^{ab}_{lh} > S^{ab}_{hl} > S^{ab}_{ll} > 0$. This different ranking of surpluses would change Proposition 1: in this case, in Period 2, people would marry a type with their exact traits, rather than a type with the same number of high traits. Conversely, all the other propositions do not change. In particular, there would still be assortative matching in attractiveness in Period 1. To see why this is the case, consider the case in which all healthy people marry in Period 2. In that case, there is assortative matching in attractiveness in Period 1 because attractive women, who are all unhealthy, maximize marriage surplus by marrying attractive men, who are also all unhealthy, rather than marrying unattractive and unhealthy men, which is the other alternative option. Suppose now that some healthy women prefer to marry in Period 1. There is assortative matching in attractiveness also in this case. If the average surplus from marrying an attractive partner is higher than marrying an unattractive partner, then all want to marry attractive men, but attractive men prefer to marry attractive women. On the other hand, if the average surplus from marrying an attractive partner is lower than marrying an unattractive partner, then all want to marry unattractive men, but unattractive men prefer to marry unattractive women. Propositions 3 to 5 are also unchanged qualitatively, as they do not depend on the different ranking of surpluses. However, the incentives for delaying marriage are now higher for healthy women, as the cost of marrying an unhealthy partner in Period 1 is greater than in the previous version of the model. Therefore, making marrying unhealthy men less appealing exacerbates the adverse selection in the marriage market in Period 1.

In this setting, making health observable in Period 1 does not change the signs of the predictions. If anything, it makes the Period 1 increases in marriage rates larger for healthy women, as more of them delayed marriage in Period 2 when health is unobserved. This version of the model also implies that making health observable in Period 1 has no effect on the marital status and surplus of already married women, since divorce is ruled out and the average surplus does not change.

6.2 Divorce Costs

We modeled divorce costs as being sufficiently high to prevent married people from divorcing. If we relax this assumption and, for simplicity, make divorce costless, Propositions 1 to 3 do not change, while 4 and 5 do. The main difference is that everybody marries a person with the same attractiveness in Period 1 because doing that gives a strictly positive surplus. In Period 2, people who are mismatched divorce and marry a partner with the same number of high traits.

In this case, making health observable in Period 1 does not change the likelihood of being married in Period 1 (as everybody marries in Period 1 anyway). Conversely, as before, making health observable increases Period 1 surplus for healthy women and decreases Period 1 surplus for unhealthy women. This occurs because now all women marry a partner with the same number of high traits, while, when health is unobserved, healthy women marry partners with fewer high traits on average, and unhealthy women marry partners with more high traits on average.

Making health observable in Period 1 also increases divorce rates for women already married to a spouse with mismatched traits in Period 1. Those women would have divorced in Period 2, after finding out their spouses' health, but now do so In Period 1. If re-marriage is instantaneous, then we also expect the same effects on surplus for married women as the ones described for singles. That is, surplus increases for healthy women and decreases for unhealthy ones through divorce and re-marriage with a partner who has the same number of high traits.

6.3 Dependence Between Attractiveness and Patience

To simplify our notation, we assumed that patience and attractiveness are independent of each other. If we assume that attractive people are more patient, our Propositions do not change. In fact, now a higher fraction of attractive than unattractive women wait to marry in Period 2, and, therefore, making health observable in Period 1 causes an even bigger increase in the likelihood of marrying in Period 1 for attractive than unattractive women. All the other predictions are unchanged.

6.4 Dependence Between Traits and Gender

To simplify our notation, we also assumed that attractiveness and health are independent of gender. Making this assumption results in a Period 2 equilibrium in which each woman is matched with a partner with the same number of high traits, because each type has equal size for men and women. If we relax this assumption and, for example, have a higher proportion of unhealthy women than men, the spirit of the propositions does not change - we still have positive assortative matching in the number of high traits in Period 2 and in attractiveness in Period 1, and sufficiently patient healthy people who wait to marry in Period 2. However, in this case there are two differences. First, some women 'marry down,' that is, marry a man with fewer high traits in Period 2 and marry a man of lower attractiveness in Period 1. Second, some unhealthy women may remain unmarried in Period 1. This is because there are more men than women who want to wait and marry in Period 2. Something similar would also occur if men are more patient than women. In this setting, making health observable in Period 1 increases the marriage likelihood in Period 1 also for unhealthy women and may increase or decrease their Period 1 surplus.

6.5 Own Health is Unobservable

The last case we consider is a scenario in which people do not know their own health in Period 1. In that case, Propositions 1 and 2 do not change: marriage is still assortative in the number of high traits in Period 2 and in attractiveness in Period 1. Propositions 3 and 4 change because all people have an expected health equal to the population average and will all behave identically, conditional on attractiveness. Therefore, sufficiently patient healthy and unhealthy women marry in Period 2, while impatient healthy and unhealthy women marry in Period 1. Proposition 5 does not change as, under some conditions, attractive women are more likely to marry in Period 2 than unattractive women.

If health becomes observable in Period 1, marriage likelihood increases for both healthy and unhealthy unmarried women, as patient women no longer have to wait, and, possibly, more for attractive than unattractive women. The average surplus increases for healthy women because (i) all, rather than some, marry in Period 1 and (ii) none of them marries an unhealthy man. Conversely, the effect of making health observable in period 1 on unhealthy women's surplus is unclear. This occurs because patient unhealthy women now marry in Period 1, generating a positive marriage surplus. However, the average surplus for impatient unhealthy women decreases, as now they all marry unhealthy men, while, with health unobserved, some of them would have married healthy men.

References

Baird, Sarah, Erick Gong, Craig McIntosh, and Berk Ozler, "The Heterogeneous Effects of HIV Testing," *Journal of Health Economics*, September 2014, *37*, 98–112.

Beegle, Kathleen, Michelle Poulin, and Gil Shapira, "HIV Testing, Behavior Change, and the Transition to Adulthood," *Economic Development and Cultural Change*, July 2015, *63* (4), 665–684.

Berendes, Sima and Rajiv Rimal, "Addressing the Slow Uptake of HIV Testing in Malawi: The Role of Stigma, Self-efficacy, and Knowledge in the Malawi BRIDGE Project," *Journal of the Association of Nurses in AIDS Care*, May/June 2011, *22* (3), 215–228.

Berge, Erling, Daimon Kambewa, Alister Munthali, and Henrick Wiig, "Lineage and land reforms in Malawi: Do matrilineal and patrilineal landholding systems represent a problem for land reforms in Malawi?," *Land Use Policy*, November 2014, *41*, 61–69.

Bergstrom, Theodore and Mark Bagnoli, "Courtship as a Waiting Game," *Journal of Political Economy*, February 1993, *101* (1), 185–202.

Bongaarts, John, "Late marriage and the HIV epidemic in sub-Saharan Africa," *Population Studies*, 2007, *61* (1), 73–83.

Boozer, Michael and Tomas Philipson, "The Impact of Public Testing for Human Immunodeficiency Virus," *Journal of Human Resources*, Summer 2000, *35* (3), 419–446.

Chesney, Margaret and Ashley Smith, "Critical Delays in HIV Testing and Care: the Potential Role of Stigma," *American Behavioral Scientist*, April 1999, *42* (7), 1162–1174.

Delavande, Adeline and Hans-Peter Kohler, "The Impact of HIV Testing on Subjective Expectations and Risky Behaviors in Malawi," *Demography*, August 2012, *49* (3), 1011–36.

Denison, Julie, Kevin O'Reilly, George Schmid, Caitlin Kennedy, and Michael Sweat, "HIV Voluntary Counseling and Testing and Behavioral Risk Reduction in Developing Countries: A Meta-analysis, 1990-2005," *AIDS and Behavior*, May 2008, *12* (3), 363–373.

Gong, Erick, "HIV Testing and Risky Sexual Behavior," *Economic Journal*, February 2015, *125* (582), 32–60.

Hainmueller, Jens, "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies," *Political Analysis*, 2012, *20*, 25–45.

—— **and Yiqing Xu**, "ebalance: A Stata Package for Entropy Balancing," *Journal of Statistical Software*, 2013, *54* (7).

Harwood-Lejeune, Audrey, "Rising Age at Marriage and Fertility in Southern and Eastern Africa," *European Journal of Population*, September 2001, *17* (3), 261–280.

Hosegood, Victoria, Nuala McGrath, and Tom Moultrie, "Dispensing with marriage: Marital and partnership trends in rural KwaZulu-Natal, South Africa 2000-2006," *Demographic Research*, 2009, *20*, 279–312.

Kaler, Amy, ""Many Divorces and Many Spinsters": Marriage as an Invented Tradition in Southern Malawi, 1946-1999," *Journal of Family History*, October 2001, *26* (4), 529–556.

—— , "'When They See Money, They Think it's Life': Money, Modernity, and Morality in Two Sites in Rural Malawi," *Journal of Southern African Societies*, June 2006, *32* (2), 335–349.

Katz, Lawrence, Jeffrey Kling, and Jeffrey Liebman, "Moving to Opportunity in Boston: Early Results of a Randomized Mobility Experiment," *Quarterly Journal of Economics*, May 2001, *116* (2), 607–654.

Magruder, Jeremy, "Marital Shopping and Epidemic AIDS," *Demography*, November 2011, *48* (4), 1401–1428.

Marston, M, E Slaymaker, S Floyd, N McGrath, I Kasamba, T Lutalo, M Nyirenda, A Ndyanabo, Z Mupambireyi, and B Zaba, "Trends in marriage and time spent single in sub-Saharan Africa: a comparative analysis of six population-based cohort studies and nine Demographic and Health Surveys," *Sexually Transmitted Infections*, 2009, *85*, i64–i71.

Matovu, Joseph and Fredrick Makumbi, "Expanding access to voluntary HIV counselling and testing in sub-Saharan Africa: alternative approaches for improving uptake, 2001-2007," *Tropical Medicine and International Health*, November 2007, *12* (11), 1315–1322.

Maughan-Brown, Brendan and Laura Nyblade, "Different Dimensions of HIV-Related Stigma May Have Opposite Effects on HIV Testing: Evidence Among Young Men and Women in South Africa," *AIDS and Behavior*, May 2014, *18* (5), 958–965.

Meekers, Dominique, "The Process of Marriage in African Societies: A Multiple Indicator Approach," *Population and Development Review*, March 1992, *18* (1), 61–78.

Mukiza-Gapere, Jackson and James Ntozi, "Impact of AIDS on Marriage Patterns, Customs, and Practices in Uganda," *Health Transition Review*, 1995, *5S*, 201–208.

Pinto, Andrew, Monique van Lettow, Beth Rachlis, Adrienne Chan, and Sumeet Sodhi, "Patient Costs Associated with Accessing HIV/AIDS Care in Malawi," *Journal of the International AIDS Society*, March 2013, *16*, 1–6.

Reniers, Georges, "Marital Strategies for Regulating Exposure to HIV," *Demography*, May 2008, *45* (2), 417–438.

Sambisa, William, Sian Curtis, and Vinod Mishra, "AIDS stigma as an obstacle to uptake of HIV testing: evidence from a Zimbabwean national population-based survey," *AIDS Care*, February 2010, *22* (2), 170–186.

Schatz, Enid, "'Take your mat and go!': Rural Malawian women's strategies in the HIV/AIDS era," *Culture, Health, and Sexuality*, Sept/Oct 2005, *7* (5), 479–92.

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

Watkins, Susan, "Navigating the AIDS Epidemic in Rural Malawi," *Population and Development Review*, December 2004, *30* (4), 673–705.

World Health Organization, "Guidance on Provider-Initiated HIV Testing and Counseling in Health Facilities," 2007.

Young, Sean and Yuda Zhu, "Behavioral Evidence of HIV Testing Stigma," *AIDS and Behavior*, April 2012, *16* (3), 736–740.