Scared Straight or Scared to Death? Fatalism in Response to Disease Risks

Jason T. Kerwin*

February 5, 2025

Click here for the latest version of this paper

Abstract

This paper shows that responses to disease risks can be "fatalistic": higher risk beliefs can lead to *more* risk-taking rather than less. Intuitively, this can occur because high risk beliefs raise not only the chance of contracting the disease (which raises the marginal cost of risk-taking) but also the perceived chance that you are already infected (which lowers the marginal cost). I test for fatalism by randomly providing information about the true (low) average risk of HIV transmission in Malawi. Consistent with rational fatalism, the treatment causes sexual activity to rise slightly on average but decline sharply for people with high initial risk beliefs—especially those with high baseline levels of sexual activity.

JEL Codes: D81, D83, I12

Keywords: Risk and Uncertainty, Health Behaviors, Health Economics, Economic Epidemiology

*Kerwin: Department of Economics and Center for the Study of Demography and Ecology, University of Washington, 342 Savery Hall, 410 Spokane Lane, Seattle, WA 98105, IZA, and J-PAL (email: jkerwin@uw.edu). This paper is dedicated to Charles Theodore "Ted" Kerwin. A previous version of this paper circulated under the title "Scared Straight or Scared to Death? The Effect of Risk Beliefs on Risky Behaviors". I am indebted to the guidance of Rebecca Thornton, Jeff Smith, John DiNardo, and David Lam. I thank Martha Bailey, Victoria Baranov, Marc Bellemare, Daniel Bennett, Yaniv Brandvain, Lasse Brune, Julie Buhl-Wiggers, Eric Chyn, Laura Derksen, Audrey Dorélien, Enda Hargaden, Anne Fitzpatrick, Willa Friedman, Paul Glewwe, Joe Golden, Erick Gong, Alan Griffith, Rachel Heath, Janna Johnson, Caroline Krafft, Anita Mukherjee, Matthew Rabin, Joe Ritter, Manisha Shah, Mel Stephens, Susan Watkins, Josef Woldense, and seminar participants at Michigan, PAA, MIEDC, ABCDE, NEUDC, IFPRI, and UMN for their helpful feedback and comments on the paper. In addition, I also thank the editor Steffen Huck and three anonymous referees for detailed comments and feedback. This project would not have been possible without the excellent fieldwork supervision efforts of Anderson Moyo, Synab Njerenga, and Christopher Nyirenda, the hard work of a team of fantastic survey interviewers, and the time and participation of the people who took part in the study. Data collection for this project was supported by grants from the University of Michigan and the Russell Sage Foundation. This research was supported in part by an NIA training grant to the Population Studies Center at the University of Michigan (T32 AG000221), as well as by fellowship funding from the Rackham Graduate School. This study was reviewed and approved by IRBs at the University of Malawi College of Medicine (COMREC, protocol # P.07/11/1107) and the University of Michigan (IRB-HSBS, protocol # HUM00052708). This study is registered with the AEA RCT Registry under registration number AEARCTR-0006666. All errors and omissions are my own.

This paper presents empirical evidence that responses to disease risks can be "fatalistic"—that higher risks can lead to more risk-taking rather than less. This contrasts with conventional "risk compensation", under which higher risks induce safer behavior (Peltzman 1975). In line with this prediction, extensive empirical research has documented a negative relationship between risks and risk-taking.[1] However, a number of studies have shown that, in theory, fatalism can be a rational response to irreversible disease risks (e.g. Philipson and Posner 1993; Kremer 1996). Intuitively, if you have already been exposed to a disease and you learn the risk of contracting the disease from each exposure is very high, this can cause you to infer that you are infected already. This drives the perceived marginal cost of additional risk-taking down to zero, leading you to take more risks.[2]

My data comes from a randomized field experiment that I designed to test for the presence of fatalism in southern Malawi. The treatment group received information about the average risk of HIV transmission from unprotected sex, which is far lower than their typical *ex ante* beliefs. Standard risk compensation predicts that this treatment, which lowers the perceived riskiness of each sex act, should lead to increased sexual activity. Fatalism, in contrast, predicts that people above a threshold level of initial risk beliefs will have less risky sex. This happens because the treatment makes fatalistic people less certain that they are doomed to HIV infection.

I analyze the data by examining how the effects of the information treatment vary by respondents' baseline risk beliefs. My analytic approach is based on a previous working paper (Kerwin 2012) which found suggestive evidence of fatalistic behavior in another part of southern Malawi.[3] To tie my hands in the analysis of my data, I use the same survey questions to capture risk beliefs and sexual activity as in that study, and define the variables I use in the same way. This leaves the choice of the functional form of the variation of treatment effects as the key remaining degree of freedom. To address this, I show that I find

[1] See e.g. Oster (2012) on HIV, Viscusi (1990) on smoking, and Gayer, Hamilton, and Viscusi (2002) on hazardous waste.

[2] This theoretical mechanism is the within-disease version of the Dow, Philipson, and Sala-i Martin (1999) model of competing health risks. It is also similar to the Oster (2012) finding that increased mortality risks from other diseases can lead to smaller behavioral responses to HIV.

[3] This paper is still publicly available on UC Berkeley's website, and on the Internet Archive's Wayback Machine: https://web.archive.org/web/20130801000000*/https://cega.berkeley.edu/assets/cega_events/4/Kerwin-Jason_Rational-Fatalism_2012.pdf.

the same pattern of results whether I examine treatment effects by decile or instead use a non-parametric approach that uses local linear regressions. Moreover, my results are robust to running my regressions without any control variables.

My results confirm that people with high initial risk beliefs respond fatalistically to HIV risks: the information treatment causes them to have less sex rather than more. The treatment has positive or null effects on sexual activity for most people, but causes a statistically significant decline of nearly 50% for the top decile of baseline risk beliefs. Moreover, I find that fatalism is stronger for people who have higher levels of baseline sexual activity (and hence more past exposures to HIV). This is consistent with fatalism being driven more by past exposures to HIV rather than inevitable future exposures, but does not rule out either channel.[4] I also find that the information treatment led to higher rates of self-reported HIV testing among fatalistic people—shifting them from having below-average to above-average testing rates.

This paper contributes to a growing empirical literature that studies how people's subjective risk beliefs affect their behavior, building on the foundational work of Manski (2004). Recent research has shown that it is possible to collect subjective risk beliefs in developing-country settings (Delavande, Giné, and McKenzie 2011). Moreover, subjective risk beliefs also drive behavior in domains ranging from HIV (e.g. Dupas 2011) to water safety (e.g. Bennear et al. 2013) to migration (e.g. Shrestha 2020).

The possibility of fatalism as a response to HIV is well-known theoretically. The results in this study add new empirical evidence to a rich existing theoretical literature on the topic of fatalism in response to risks. The first study to establish the possibility of fatalistic responses to disease risks was Philipson and Posner (1993), who point out that if the prevalence of a disease becomes sufficiently high then people may have "nothing more to lose" (p. 49). Kremer (1996) introduces the term "fatalistic" for this phenomenon, and shows that it can lead to multiple steady-state equilibria. O'Donoghue and Rabin (2001) note that the cost function for HIV exposures is concave and thus has a tipping point where the marginal cost falls rather than rising with increased risks. Matthies and Toxvaerd (2023) show a closely

[4] One way of determining whether fatalism is solely driven by past exposures rather than inevitable future ones would be to test some members of the treatment group for HIV. My study does not do this.

related result: increases in risk aversion can lead people to seek out more, rather than less, disease exposure. They find that this result can exacerbate rational fatalism. Related work by Shapiro and Wu (2011) studies a different notion of fatalism: in their model, fatalism means misperceiving the return to effort to be lower than it actually is.

I build on this existing theoretical literature by providing empirical evidence of fatalism. Despite the extensive body of theoretical work on fatalism, there is only limited empirical evidence of its existence. Ethnographic work has documented fatalistic reasoning about HIV in Malawi (Kaler 2003, Kaler and Watkins 2010) and Uganda (Barnett and Blaikie 1992), while Matthies and Toxvaerd (2023) provide evidence in support of their model from a lab experiment. In previous research on another part of southern Malawi (Kerwin 2012), I show that the cross-sectional relationship between sexual behavior and the perceived HIV infection risks from unprotected sex, rather than being downward-sloping, is U-shaped—which is consistent with fatalistic behavior. I also develop a theoretical framework based on O'Donoghue and Rabin (2001) to explain this pattern of fatalistic behavior. Wilson, Xiong, and Mattson (2014) cite a reduction in fatalism as the likely mechanism for their finding that circumcision (which protects against HIV transmission) leads people to have less unprotected sex, rather than more.[5] However, their data does not allow them to determine whether the intervention shifted people's perceived risks.[6] There is also evidence that people who are told they are HIV-positive have more unprotected sex (Gong 2015), which is consistent with the mechanism behind fatalism, but does not demonstrate that exaggerated beliefs about the riskiness of unprotected sex cause fatalistic behavior.

My study advances this literature by demonstrating the empirical pattern predicted by models of rationally fatalistic responses to disease risks: people with low and moderate initial risk beliefs exhibit typical risk compensation, while those with high risk beliefs are fatalistic. The existence of fatalism is an important consideration for designing policy responses to the HIV pandemic. While standard epidemiological models do not allow for behavior changes in

[5] A related phenomenon is documented in Baranov and Kohler (2018), who show that lowering the mortality risk from HIV leads to higher savings and human capital investments. This is distinct from fatalism: it is a response to a lower cost of HIV conditional on contracting it, rather people feeling they are doomed to get HIV no matter what.

[6] Moreover, Godlonton, Munthali, and Thornton (2016) find no effects on sexual behavior when they provide circumcised men with information about the HIV transmission benefits of circumcision.

response to disease risks, Kremer (1996) and Greenwood et al. (2019) show that behavioral responses to disease risks have crucial effects on the spread of HIV. The average effect of the specific information treatment in my study is to increase sexual activity, although this effect is fairly small. The optimal design of an information campaign about HIV risks will depend on the relative importance of fatalistic and non-fatalistic groups in driving the spread of the disease. Some epidemiological models suggest that small groups of high-activity people such as sex workers and their clients can play a disproportionate role in the spread of the virus (see e.g. Koopman, Simon, and Riolo 2005). The key factors here are unobservable in my study: the role of fatalistic people in the overall epidemic depends crucially on how much sex they have with each other and how they are connected to other people through the sexual network.

Policymakers should consider the possibility of fatalism in response to other disease risks as well. The same basic mechanism behind the fatalistic responses to HIV may hold for any condition perceived to be binary and where one's disease status cannot be observed immediately.[7] Potential examples include STIs such as HSV-2, exposure to cancer-causing chemicals—and also COVID-19. Given that people may react fatalistically to these risks, it is not clear that "scared straight" style information campaigns that exaggerate the transmission rates of diseases will reduce risk-taking. Conversely, it is also unclear whether telling people the true risks is optimal from a disease control perspective; this depends on the relative importance of different groups in driving the overall epidemic. It also depends on the general equilibrium effects of providing risk information: changes in the demand for risky sex by high risk groups may affect the willingness of other groups to have sex.

Anecdotal evidence suggests that many people perceive COVID-19 to be so contagious that contracting it is inevitable, and commentators have raised concerns about people becoming fatalistic in response to the pandemic (e.g. Cowen 2020, Oster 2020). A recent survey experiment echoes those concerns: Akesson et al. (2022) show that increasing people's perceptions of the contagiousness of COVID-19 reduces their stated willingness to engage in

[7] Indeed, cross-sectional evidence also suggests that fatalism may affect decisions about dieting, preventive healthcare, and smoking cessation (Ferrer and Klein 2015). Similarly, Adda (2007) find that responses to Mad Cow disease vary by past exposure in a pattern that is consistent with fatalism, although this is not explicitly discussed in the paper.

social distancing, which is exactly in line with the predictions of rational fatalism and my empirical results. I go beyond the average effects documented by Akesson et al., showing that the pattern of treatment effect heterogeneity by baseline risk beliefs matches what fatalism models predict. My results show that these sorts of perverse effects could indeed be explained by fatalism, and demonstrate that fatalism has consequences for real-world behavior. They therefore suggest we should be cautious in the use of "scared straight"-style messaging about disease risks. Emphasizing that an activity's risks are extremely high—in the name of encouraging safer behavior—can backfire, causing people to fatalistically take even more risks.

1 Experiment and Data

My data comes from a randomized field experiment I conducted in the Zomba district of southern Malawi from August to December of 2012. I chose this location because both ethnographic work on southern Malawi (Kaler 2003, Kaler and Watkins 2010) and my own previous quantitative research on Zomba district (Kerwin 2012) are suggestive of fatalism. I randomly selected 70 villages from one sub-district, assigning half to the control group and half to the treatment group. I then randomly sampled 30 adults aged 18-49 from each village. The village sample was stratified by distance to the nearest trading center; within each village, the sample of adults was stratified by gender. The baseline survey excluded sexually inactive people because they are unlikely to be fatalistic. This exclusion was imposed during the baseline survey by skipping people who had never had sex to the end of the survey; it removed just 40 observations (2.6% of the sample). See Appendix B for further details about the sampling procedure.

This sampling and exclusion process yielded a final sample of 1,503 completed baseline surveys. Interviewers re-contacted the original respondents for an endline survey approximately 1-4 months later, successfully locating 1,292 of them. The random assignment produced study arms that were fairly balanced on baseline covariates. Table 1 shows treatment-control balance tests separately for the non-fatalistic (bottom nine deciles of baseline risk

beliefs) and fatalistic (top decile) subsamples of my data.[8] Out of 44 total tests, 4 of them (9%) reject null hypothesis at the 5% level. This is more than we would expect under the null if the tests were independent, but they are correlated with one another. There is some evidence of imbalance on sexual activity variables for the fatalistic subsample. To address this, my main specification controls for all the variables in this table and their interactions with the treatment indicator. Attrition rates were balanced across study arms, and there is no evidence of differential attrition by baseline characteristics (Appendix Tables A2 and A3).

The treatment was an information script, presented at the end of the baseline survey, that told respondents the average annual risk of HIV transmission from an infected to an uninfected sex partner who are having unprotected sex about once every three days. I chose the annual risk because it is simpler to explain than the per-act risk (which is very small), and also because it is available on the Malawi National AIDS Commission's website (Malawi National AIDS Commission 2009, p. 11). The average risk is about 10% per year (Wawer et al. 2005); this corresponds to a per-sex-act risk of approximately 1 in 1000, since couples in the Wawer et al. study had sex about 100 times per year on average.[9] The risk of HIV transmission varies around this average: for example, it is over 50% lower for circumcised men (Gray et al. 2007), and up to 96% lower if the infected partner is on antiretroviral treatment (Cohen et al. 2011). However, these variations in the risk are dwarfed by the miscalibration of people's priors; the median person in the sample overstates the per-act transmission risk by a factor of nearly a thousand. The intervention received ethical approval from IRBs at the University of Malawi College of Medicine and the University of Michigan. For a discussion of the ethical dimensions of teaching people the true average risk of HIV transmission, see Appendix C. For further details about the information treatment, see Appendix D.

My presumption was that respondents would update their other beliefs about both the transmission rate and also the prevalence of HIV in response to this information, since the

[8] For the balance table for the full sample, see Appendix Table A1.

[9] Wawer et al. do not directly calculate an overall annual transmission rate. They report 68 total seroconversions (page 1) and a total of 553 ten-month followup periods (Table 2) for an overall annual transmission rate of 14.7%. I chose the rate of 10% as it is the upper bound of the range reported by the Malawi National AIDS Commission, and because it is easier to explain; I presented the risks as annual risks to ease explanations as well.

prevalence is a function of past transmission. Table A4 shows that this did in fact happen. The treatment reduces the perceived annual risk from unprotected sex by 38 percentage points (relative to a control-group mean of 91%), and the per-act risk by 37 percentage points (control mean = 74%). Respondents also update their beliefs about condom-protected risks and prevalence.

There are two potential limitations to this information treatment. First, people may not feel that average risks apply to them directly, instead feeling that their own risk is lower (Weinstein 1989). Even if this is true, the average risks measured in the survey and targeted by the information treatment should drive updating in people's perceived personal risks (similar to the aforementioned updating of prevalence beliefs). Consistent with this, I find that the information treatment leads to changes in actual behaviors. Second, the information treatment might have changed the perceived costs of HIV or made them more salient. There is no evidence of this: Appendix Figure A5 shows that there are no average treatment effects on how long people think they will live if they contract HIV, and no meaningful treatment effect heterogeneity by baseline beliefs either.

To minimize the chance of contaminating the control villages, the treatment-group baseline surveys were done after the control-group baseline surveys were completed, following Godlonton, Munthali, and Thornton (2016). For the same reason, the survey interviewers were trained to administer the information intervention only after the end of the baseline survey wave for the control group.[10]

Since the goal of the experiment was to test for fatalism, I considered stratifying the randomized treatment assignment by levels of baseline risk beliefs. Doing so would have increased my statistical power, but was infeasible because of other elements of the design. In particular, since the information treatment was randomized at the village level, I could not target it solely at some of the sample. I would also have had to conduct much of the baseline survey to collect the risk beliefs needed to do the stratification, greatly increasing costs.

[10] Kerwin and Ordaz Reynoso (2021) directly test for spillovers, finding no evidence that people with more treated social ties have different HIV risk beliefs.

1.1 Measures of Sexual Activity

I conducted this experiment to test the preliminary empirical evidence of fatalism from Kerwin (2012). In that study I also pilot-tested the survey questions that I used in the current study. I therefore use the same primary outcome variable as I used in that earlier paper: the total number of sex acts recorded on a retrospective "diary", in which the interviewer walked respondents through the previous seven days. I also examine the robustness of my findings to other outcomes. These include simple recall questions about sex acts or sex partners in the past 30 days, as well as a combined outcome index that uses all the sexual activity variables from the survey. I construct the index by taking the first principal component of the control-group data and applying those weights to the treatment group as well; I do this separately for the baseline and endline waves. This approach is based on the assumption that each outcome is a valid measure of the same underlying construct (sexual risk-taking).

I do not use STI incidence to measure risky sex for two reasons. First, my project budget would not have allowed me to collect biomarkers. Second, it is not clear that any STI biomarkers would be useful for detecting fatalism. The only common STI in southern Malawi (other than HIV) is HSV-2 (Baird et al. 2012). Since HSV-2 is not curable, using it as my outcome measure would screen out many of the highest-risk individuals who are the focus of my study: many would already have HSV-2 at baseline, and thus be coded as a 1 for this outcome variable at endline irrespective of the effect of the information treatment. I do collect data on condom purchases, which I offered at the end of the endline survey, following Thornton (2008). However, 80% of my sample reports that condoms are available for free, so I do not focus on this as an outcome measure; I do include it in the combined outcome index. Another potential outcome is pregnancy, but only 12 women (1.6% of my sample) conceived a child between the baseline and endline survey waves, so I have no statistical power for this outcome variable.

Self-reports of sexual activity may be affected by response biases and thus mismeasure actual risky sex (Johnson and Delamater 1976). There are three reasons to think that this should have minimal effects on my findings. First, STI biomarkers also measure risky sex with error: the tests have non-zero false-positive and false-negative rates; in addition, not every

risky sex act leads to an STI infection. Given the latter issue, Corno and De Paula (2015) argue that self-reported sex is a better proxy for risky sexual behavior than STI biomarkers when STI prevalence is low, as is true in Malawi for curable STIs. Second, my method of measuring sexual activity was extensively validated and is designed to build rapport and comfort for respondents. I use a retrospective "diary"-based approach to measuring sexual behavior that I validated through previous work on sexual behavior in southern Malawi (Kerwin 2012). This approach builds on research that shows that calendar-based methods reduce recall bias compared with single-question recall methods (Belli, Shay, and Stafford 2001, Luke, Clark, and Zulu 2011).

Third, to create the qualitative pattern of results that I find in the paper, any response biases would have to be different for the treatment and control groups, but in a way that reverses itself as risk beliefs rise. The response biases could plausibly differ across study arms on average. The information treatment did not provide any guidance on sexual activity, and contained no judgmental language about how much sex people should have, but it did say how much sex people in the Wawer et al. (2005) study were having. This could conceivably provide a role model effect (although the subjects of that study were from a different country). However, even interventions that directly try to target sexual behavior change have average effects that are very small (Oster 2012). Moreover, this information should not have affected respondents differently by the combination of treatment status and baseline risk beliefs.

1.2 Measures of Risk Beliefs

To measure subjective beliefs about HIV risks I use a set of questions about proportions out of a fixed number of people. For example, one of the questions is "If 100 men, who do not have HIV, each sleep with a woman who is HIV-positive tonight and do not use a condom, how many of them do you think will have HIV after the night?"[11] I then divide the total number by the denominator to produce a probability. I collected transmission risk beliefs for both per-act risks and annual risks, and for sex with and without condoms.

[11] This is the male version of the question; the genders are inverted for the female version. See Appendix Figure A1 for the exact phrasing of the questions used. Following Hudomiet, Kézdi, and Willis (2011), I ask respondents who report beliefs of 50% whether they were just not sure, and replace their initial answer with their best guess if so.

People begin with extremely high risk beliefs: the median control-group respondent believes that a single unprotected sex act with a randomly chosen sex partner has a 4 in 10 chance of giving them HIV. These exaggerated risk perceptions are consistent with what students are taught in schools in Malawi. The textbooks for the course that covers HIV prevention in secondary school (Life Skills) reference the transmission rate only once. Page 61 of Kadyoma et al. (2012) describes a young woman who contracted HIV the first time she had sex, implying a transmission rate of 100%.

My main belief variable, which I refer to as "risk beliefs" in the remainder of the paper, is a measure of the riskiness of a sex act with a randomly chosen sex partner. I construct this by multiplying the per-act transmission rate belief by the perceived prevalence of HIV among attractive people. I take this approach because there is a large mass point at 100% in the transmission rate belief distribution (Appendix Figure A2, Panel A). Since many of those people believe the prevalence of the virus is very low, their effective risk from unprotected sex is not high. Multiplying the transmission rate by the prevalence focuses on the actual risk that a given sex act carries. I discuss this choice further in Appendix G.6.1, and present results separately for each risk belief component. I use the prevalence among people the respondent finds attractive, rather than among all members of the opposite sex, to focus on the risk from potential sex partners—which may differ from the rest of the population.[12] This composite variable is the same definition of risk beliefs that I used in previous work that found suggestive evidence for fatalism in southern Malawi (Kerwin 2012).

A separate paper using the same dataset (Kerwin and Ordaz Reynoso 2021) documents that there is a small but detectable interviewer effect on recorded baseline risk beliefs. This happened because the interviewers delivered the information treatment, and they did not learn the information themselves until after finishing the control-group baseline surveys. This led them to record slightly lower risk beliefs at baseline for the treatment group, likely due to subtle changes in the exact delivery of the survey questions; the evidence suggests this did not affect actual beliefs. My results are robust to correcting for this issue (Appendix G.6.3).

[12] The set of attractive people may differ from the set of prospective sex partners. I find qualitatively similar results if I instead use the probability that the respondent's primary sex partner is HIV-positive, for respondents who are in committed relationships (see Appendix G.6.4.)

2 Empirical Strategy

My main analyses focus on how the effects of the information treatment vary by people's baseline risk beliefs. I estimate the following regression:

$$y_i = \sum_{k=1}^{10} \left[\alpha_k r_i^k + \beta_k T_i \times r_i^k\right] + \boldsymbol{X_i'\gamma} + \varepsilon_i \tag{1}$$

where y_i is the outcome, the r_i^ks are indicators for quantiles of baseline risk beliefs, T_i is the treatment indicator, $\boldsymbol{X_i}$ is a vector of controls, and ε_i is a mean-zero error term. I log the outcome variable, using the Ravallion (2017) transformation to allow for zeroes and negative values.[13]

The coefficients β_k give the treatment effect for each quantile of baseline risk beliefs. In my main specification, $K = 10$, so these are deciles. There is no omitted category for r_i^k and also no main effect for the treatment; instead, I estimate all 10 decile-specific treatment effects. Similarly, there is no intercept in Equation 1 because I include main effects for all K quantiles of baseline risk beliefs.

The average effect of the treatment on the outcome is identified because the random assignment of the treatment guarantees that $E[\varepsilon_i|T_i] = 0$. However, the heterogeneity in treatment effects captured by the coefficient on $T_i \times r_i^k$ is still subject to potential omitted-variable bias. Since baseline risk beliefs are not randomly assigned, apparent variation in treatment effects by r_i^k could actually be due to other factors that are correlated with r_i^k, which themselves cause treatment effect heterogeneity. This is an important concern because risk beliefs are positively correlated with sexual behavior (see Panel B of Appendix Table F2). Moreover, fatalistic people (those in the top decile of baseline risk beliefs) differ from the rest of the sample in several ways: they have been sexually active for longer, are older, and are more likely to be Christian than people in the bottom nine deciles of risk beliefs (Table A1, Panel B).

To address this omitted-variable bias, as well as any potential baseline imbalance, my controls $\boldsymbol{X_i}$ include both main effects and interactions with the treatment for an extensive

[13] The Ravallion (2017) transformation is given by $h(y) = \begin{cases} 0.5(e^y - e^{-y}) - \ln(2) & \text{if } y \leq 0 \\ \ln(y + \sqrt{y^2+1}) - \ln(2) & \text{if } y \geq 0 \end{cases}$

set of exogenous covariates. These include the baseline value of the outcome variable,[14] as well as all the baseline covariates from Table 1. Following Imbens and Rubin (2015), I demean all these covariates prior to constructing the interaction terms, so the main effects of T_i and $T_i \times x_i$ can still be interpreted as the sample-average effects. In my robustness checks I show that my results are not sensitive to the inclusion of any of the control variables. I also control for sampling strata indicators to improve efficiency (Bruhn and McKenzie 2009).

I cluster the standard errors by village, which is the level at which the treatment was randomized. To address multiple hypothesis testing, I show Montiel Olea and Plagborg-Møller (2019) sup-t simultaneous confidence bands rather than pointwise confidence intervals when graphing my results. I use these simultaneous confidence bands for inference as well: unless otherwise noted, all p-values reported in the text are sup-t adjusted. These confidence bands guarantee 95% coverage for the entire function of estimates, and thus are conservative even relative to other multiple testing adjustments.

2.1 Non-parametric Estimates

The estimates of the quantile-specific treatment effects from Equation 1 depend on the choice of the number of quantiles to use. As a robustness check I examine 5, 15, and 20 quantiles in addition to 10. I also take a non-parametric approach to show how the treatment effects vary by. To do this I run local linear regressions of the outcome on baseline risk beliefs, separately by treatment status.[15] I then compute the difference between the treated and control predicted values as an estimate of the belief-specific baseline risk belief.

To do inference on these estimates I use the Bayesian bootstrap of Rubin (1981). I draw 1,000 random weights for each observation, clustered by village and stratified by stratification cell.[16] I then use these weights to estimate a distribution of estimated belief-specific treatment effects, and construct the 95% confidence interval as the range between the 2.5^{th} and 97.5^{th}

[14] Controlling for baseline values of the outcome improves precision; in Appendix E I show via Monte Carlo simulation that it also reduces finite-sample bias if there is any baseline imbalance in the outcome (irrespective of statistical significance).

[15] I use an Epanechnikov kernel and a data-driven rule-of-thumb bandwidth.

[16] I draw the weights from an exponential distribution and normalize them to sum to 1. The standard bootstrap can be viewed as drawing random integer weights that can be zero. In my setting, this means that some values of baseline beliefs would have no predicted values in certain draws; the Bayesian bootstrap avoids this issue.

percentiles of this distribution.

3 Results

My main results are shown in Figure 1. Panel A shows the effects on endline risk beliefs by the level of baseline risk beliefs. The changes in risk beliefs are larger for people with initially higher beliefs, but this relationship is somewhat noisy and may not be monotonic.[17] It is not clear whether we should expect the updating of risk beliefs to be monotonic. On one hand, people with higher risk beliefs experienced a larger shock to their priors due to the treatment, and thus should update more. On the other, under the conventional Bayesian updating model with a single prior, beliefs closer to 100% imply greater certainty and thus less updating relative to beliefs near 50%.[18]

Panel B plots the treatment effect on the y-axis against deciles of baseline risk beliefs on the x-axis. All ten deciles are displayed, positioned at the average value of the baseline risk belief for each decile. That is, the x-axis shows the average level of the risk belief within each decile of risk beliefs. Note that there is no main effect for the treatment indicator, and no omitted category of baseline risk beliefs here; the figure shows the total treatment effect for each baseline risk belief decile. Because the information treatment reduces people's risk beliefs, under fatalism we would expect negative treatment effects for people with high values of baseline risk beliefs.

The pattern of heterogeneity in treatment effects by baseline risk beliefs is highly non-linear: there are positive or zero treatment effects for the first nine deciles, and a large negative effect (67 log points) for the highest decile of beliefs. This impact, which is 49% using the transformation $\beta_{percent} = exp(\beta) - 1$, is very similar to the effect that Gong (2015). He tells people with high priors about their HIV status that they do not have HIV and finds that risky sex declines by 59%. My results are also comparable in magnitude to those from Godlonton, Munthali, and Thornton (2016), who find that uncircumcised men reduce their sexual risk-taking by 0.18 SDs in response to information about the HIV transmission

[17] Appendix Figure A3 shows histograms of endline risk beliefs by study arm.

[18] If people instead have multiple priors (Gilboa and Schmeidler 1989) then people with higher baseline risk beliefs do not necessarily update less, since they may not be as certain about their priors.

benefits of circumcision. My larger effects (equivalent to 0.34 SDs) may be because the information shock I provided is larger.

The fatalism effect for the highest level of baseline risk beliefs is strongly statistically significant ($p = 0.01$). There are small positive effects for the bottom nine deciles, with an average increase of just 15%. While the estimated effects are slightly negative at the fourth and seventh deciles, I can reject the equality of those effects with the tenth-decile effect (p=0.09 and 0.06 respectively). Appendix Table A5 shows the same results numerically. What are the most likely values of the treatment effect for the top decile of baseline risk beliefs? I can convert frequentist confidence intervals to a Bayesian credible interval if I assume that the parameter is normally distributed and impose a non-informative (i.e., uniform) prior (Bolstad and Curran, 2017, p. 242). Under these assumptions, there is a 50% chance that the treatment effect lies between -0.54 and -0.79.

In Figure 2, I show the results using the non-parametric method described in Section 2.1. The patterns are qualitatively similar to the decile-based approach. The information treatment reduces endline risk beliefs throughout the distribution of baseline beliefs, with the largest effects in the middle of the risk belief distribution. For sexual activity, we can see that the estimated treatment effect is small and positive for most people, and turns negative for people with the highest baseline risk beliefs. The tipping point into fatalistic behavior happens at a baseline risk belief of roughly 80%.

My main results are also visible in the raw data. In Appendix Table A6, I present the means and SDs of the (unlogged) outcome for each decile of beliefs at baseline and endline, as well as their difference. The same pattern of fatalism for the top decile of baseline risk beliefs is visible in the unadjusted endline data (column 6), as well as in the difference in differences (column 9), but there is no evidence of such a pattern at baseline (column 3). The fatalism effect is also visible if we look solely at changes over time within the treatment group: for the top decile of risk beliefs, weekly sex acts go down by 0.79 from the baseline to the endline (Column 8; note that the number in parentheses is the SD, not the standard error). These patterns are also visible in histograms of the raw data (Appendix Figure A4): the treatment shifts the distribution of sexual activity to the right both overall (Panel A) and for the non-fatalistic sample (Panel B) but sharply to the left for the fatalistic sample

(Panel C).

3.1 Potential Limitations

One potential limitation of my results is that the treatment effect on risk beliefs for fatalistic people is not statistically significant. Specifically, people in the top decile of baseline risk beliefs have a negative treatment effect on risk beliefs but with a 95% sup-t confidence band that crosses zero ($p = 0.14$). However, this does not mean that the treatment had no effect on beliefs for this group; statistically insignificant does not imply zero (Altman and Bland 1995). There are several reasons to believe that the treatment did shift beliefs for people at the top of the baseline risk belief distribution. First, the effect is not small: it is a 16 percentage-point reduction relative to a control-group mean of 35%. I can also characterize where the bulk of the probability mass of likely treatment effects lies using the same Bayesian assumptions I imposed above for the treatment effect on sexual activity. Under those assumptions, there is a 50% chance that the top-bracket effect on risk beliefs is somewhere between a 11 and 21 percentage-point reduction.

Second, the sup-t confidence bands are conservative for inference on pointwise effects, even relative to other multiple testing methods. In this specific case, the conventional Anderson (2008) q-value is significant at the 0.05 level ($q = 0.014$). Third, the effect for the top decile cannot be statistically distinguished from any of the other decile-specific reductions in risk beliefs. The largest reduction happens for the 7^{th} decile, at a baseline risk belief of 0.58; a t-test of the equality of the 7^{th}- and 10^{th}-decile effects has $p = 0.30$. Fourth, the non-parametric estimates in Figure 2 show that the treatment effects on beliefs are significant and negative throughout the distribution of baseline risk beliefs.

These results are robust to a wide range of other robustness checks. For the sake of space, I show these in Appendix G of the paper, but I highlight a few key ones here. First, the results do not depend on the specific number of bins that I choose for defining quantiles of baseline risk beliefs. I find similar qualitative patterns for 5, 15, and 20 bins (Appendix Figure G2). Relatedly, the negative treatment effects for people with the highest risk beliefs are also visible if I use a linear specification. In Appendix Table G1 I interact the treatment indicator with the level of baseline risk beliefs, and estimate the predicted effect for people

with a baseline risk belief of 100%. It is uniformly negative and significant at the 0.1 level in all specifications. The results are also robust to using a combined index of sexual activity as the outcome variable (Appendix Figure G1, Panel E).

3.2 Limiting Researcher Degrees of Freedom

This study did not have a pre-registered analysis plan because the AEA RCT registry did not exist at the time that it was conducted: I collected the endline data in 2012, and the RCT registry was launched in 2013 (Goyal and Cavanagh 2024). While other pre-registration options did exist, filing pre-analysis plans was not the norm in economics at the time.

Because my analyses were not pre-registered, I have taken several steps to tie my hands in the analysis in order to limit "researcher degrees of freedom" that can lead to spurious findings through specification searching (Simmons, Nelson, and Simonsohn 2011). The first is that my RCT was designed to test a specific theoretical prediction, from a publicly available working paper on the topic that was posted prior to running the experiment (Kerwin 2012).

Second, I use the exact same outcome variable (unprotected sex acts in the past week) and belief variable (perceived transmission per-act transmission rate from unprotected sex times the prevalence of HIV among attractive people in the local area) as I did in the empirical section of that earlier paper (see tables Tables 4 and 5 and Figures 4-7 of that paper). Note that I log the outcome in the main results in my current paper to simplify interpretation, but I find the same qualitative pattern if I use the unlogged outcomes (Appendix Figure G3, Panel A).

The theoretical and empirical results in my previous paper suggest that treatment effects will differ in both magnitude and sign by the level of baseline risk beliefs, but do not specify the exact functional form of this relationship. Thus the third step I have taken is to use a hands-free approach to characterizing the heterogeneity in treatment effects, using the non-parametric approach described in Section 2.1. Those results show a clear pattern of heterogeneity in treatment effects that matches the prediction of fatalism for people with the highest baseline risk beliefs.

The choice of control variables is another potential degree of freedom for analyzing data. I address this in two ways. First, my main results are robust to omitting all the controls, and to

controlling only for the sampling strata (Appendix Figure G5). Second, the non-parametric estimates include no controls.

3.3 Complementary Results

One implication of fatalism is that higher levels of baseline sexual activity should make fatalism worse; the more exposures you have already had, the more you should think you already have HIV and thus face no cost from additional risk-taking. In Figure 3 I break down my main results by whether (at baseline) people were above or below the median level of sexual activity in the past 30 days. People with below-median levels of baseline sexual activity have weaker and statistically insignificant fatalistic responses. Those above the median have much larger fatalistic responses, which are statistically significant at the 0.05 level. The existence of this pattern does not rule out a role for inevitable future sex acts in driving fatalism, since past and future sex acts are highly correlated.

Another implication of fatalism is that the reductions in sexual risk-taking should be driven by reductions in people thinking that they are doomed to get HIV no matter what. One way this could manifest is through a lower perceived chance of already having HIV. To explore this, Figure 4 shows my main results by whether people think they currently have HIV at baseline. The fatalistic responses are somewhat stronger for people who think they may be HIV-positive (Panel B) but they are statistically significant at the 10% level for both groups, and I cannot reject the equality of the top-decile treatment effects ($p = 0.459$). This result is consistent with a lack of average treatment effects on perceived HIV status (Table A4, column 9). One potential explanation for the these results is that fatalism may be driven by inevitable future HIV exposures, and thus being doomed to contract HIV in the future.

Another way the information treatment might reduce risk-taking among fatalistic people is by encouraging HIV testing.[19] In theory, testing rates should be the highest among people who are the most uncertain about their HIV status, and lower for people who think their probability of being HIV-positive is close to zero or one (Boozer and Philipson 2000). Fatal-

[19] HIV testing is commonly believed to lead to safer sexual behavior, although empirical evidence suggests the effects of a single HIV test are limited (Thornton 2008; Gong 2015).

ism makes people think their probability is nearly one; mitigating fatalism should lower this probability, making them more likely to get tested. I examine this possibility in Table 2, which presents treatment effects on self-reported HIV testing since the end of the baseline survey. Control-group respondents in the top decile of risk beliefs have much lower testing rates. The information treatment reverses that pattern, making their testing rates higher than average.

How can we reconcile the effects on self-reported HIV testing with the lack of differences in fatalism by perceived HIV status? One explanation is statistical power: out of the 112 people in the top decile of risk beliefs, just 43 think they may be HIV-positive. In contrast, the effects on testing are estimated using all 112 people in the top decile. Another possible explanation is measurement error in the HIV status variable. Malawians greatly overstate their likelihood of being HIV-positive (Anglewicz and Kohler 2009). Moreover, they update their beliefs about their HIV status very little in response to information: receiving a negative HIV test results decreases the perceived probability of being infected by less than 10 percentage points, and these effects dissipate entirely within two years (Thornton 2012) despite lasting effects on sexual behavior (Delavande and Kohler 2012). This implies that measured HIV status beliefs may diverge from the underlying drivers of people's decisions. An alternative explanation is that testing may be misreported: in Derksen, Muula, and van Oosterhout (2022), self-reported testing rates in Zomba district are over four times the rates from administrative data. However, it is unclear why exaggerated testing rates would lead to spurious treatment effects specifically for fatalistic people. Thus the pattern in Table 2 is unlikely to be entirely driven by misreporting.

4 Conclusion

I test for fatalistic responses to HIV risks by randomizing the provision of accurate information about the transmission rate of the virus—which is much lower than most people's priors. The treatment effect on sexual activity varies sharply by people's initial risk beliefs. It is slightly positive for most people but strongly negative for those with the highest initial risk beliefs. This pattern is consistent people behaving "fatalistically" in response to HIV

transmission risks, because the high risk of HIV transmission makes it likely that they will contract HIV no matter what they do.

Future research should try to uncover how people form the high risk beliefs that lead them into fatalistic behavior. Given that overestimating HIV risks seems to scare people to death, rather than scaring them straight, getting at the source of these overestimates may be crucial for understanding the continued spread of HIV in sub-Saharan Africa.

5 Tables and Figures

Table 1
Baseline Balance by Fatalism

	Panel A: Non-Fatalistic Sample				Panel B: Fatalistic Sample			
	Ctrl. Mean (SD)	Treat. Mean (SD)	Diff.† (p-value)	Obs.	Ctrl. Mean (SD)	Treat. Mean (SD)	Diff.† (p-value)	Obs.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sexual Activity								
Any Sex in Past Week	0.538	0.507	-0.032	1,163	0.609	0.512	-0.081	112
	(0.499)	(0.500)	(0.150)		(0.492)	(0.506)	(0.438)	
Total Acts in Past Week	1.800	1.613	-0.191	1,163	1.913	1.721	-0.082	112
	(2.494)	(2.385)	(0.153)		(2.331)	(2.539)	(0.864)	
Unprotected Acts in Past Week	1.559	1.461	-0.100	1,163	1.768	1.651	0.015	112
	(2.384)	(2.323)	(0.450)		(2.359)	(2.525)	(0.975)	
Sex Partners in Past 30 Days	0.804	0.804	-0.003	1,162	0.941	0.721	-0.223**	111
	(0.492)	(0.782)	(0.930)		(0.543)	(0.504)	(0.014)	
Condoms Acquired in Past 30 Days	4.454	3.538	-0.907	1,159	7.304	4.395	-3.946	112
	(14.137)	(11.619)	(0.265)		(21.148)	(12.154)	(0.407)	
Years Sexually Active	13.014	12.869	-0.127	1,152	14.147	17.293	2.235	109
	(8.276)	(8.459)	(0.794)		(8.308)	(9.644)	(0.263)	
Lifetime Sex Partners	2.981	3.586	0.583***	1,161	4.338	2.977	-1.662***	111
	(2.338)	(4.881)	(0.006)		(4.557)	(1.871)	(0.003)	
Any Chance of Having HIV	0.339	0.348	0.008	1,150	0.391	0.381	-0.017	111
	(0.474)	(0.477)	(0.793)		(0.492)	(0.492)	(0.851)	
Overall Sexual Activity Index	0.011	-0.021	-0.036	1,152	0.209	-0.104	-0.236	109
	(1.005)	(1.004)	(0.503)		(0.931)	(1.017)	(0.242)	
Demographics								
Male	0.426	0.436	0.000	1,163	0.435	0.488	0.000	112
	(0.495)	(0.496)	(0.000)		(0.499)	(0.506)	(0.000)	
Married	0.819	0.814	-0.004	1,161	0.913	0.721	-0.194**	112
	(0.385)	(0.390)	(0.861)		(0.284)	(0.454)	(0.030)	
Age	29.033	29.167	0.150	1,163	30.116	34.512	3.575*	112
	(8.423)	(8.179)	(0.754)		(8.529)	(8.738)	(0.070)	
Years of Education	5.760	5.828	0.086	1,163	5.667	6.535	0.450	112
	(3.349)	(3.445)	(0.758)		(3.328)	(3.990)	(0.583)	
Household Size	5.040	4.850	-0.197	1,163	5.029	5.186	0.138	112
	(2.223)	(2.018)	(0.234)		(2.431)	(2.185)	(0.747)	
Spending in Past 30 Days	292.497	289.436	-0.572	1,163	274.787	367.444	52.543	112
	(392.460)	(563.508)	(0.986)		(288.242)	(744.861)	(0.608)	
Assets Owned	4.149	3.941	-0.194	1,163	4.333	4.140	-0.406	112
	(2.397)	(2.296)	(0.316)		(2.571)	(2.578)	(0.411)	
Ravens Score [0-3]	1.559	1.527	-0.033	1,163	1.493	1.674	0.117	112
	(0.995)	(1.000)	(0.618)		(0.964)	(1.040)	(0.524)	
Numeracy [0-3]	0.732	0.814	0.077	1,163	0.580	0.977	0.351*	112
	(0.937)	(0.994)	(0.175)		(0.864)	(1.165)	(0.056)	
Chance of Winning Question	0.224	0.243	0.016	1,163	0.174	0.349	0.148*	112
	(0.417)	(0.429)	(0.514)		(0.382)	(0.482)	(0.057)	
Risk Attitude	0.268	0.265	-0.002	1,162	0.191	0.317	0.125	109
	(0.444)	(0.442)	(0.947)		(0.396)	(0.471)	(0.123)	
Christian	0.905	0.927	0.022	1,163	0.971	0.930	-0.053	112
	(0.293)	(0.260)	(0.388)		(0.169)	(0.258)	(0.230)	
Muslim	0.089	0.059	-0.030	1,163	0.029	0.070	0.053	112
	(0.285)	(0.236)	(0.215)		(0.169)	(0.258)	(0.230)	

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Non-fatalistic people are those in the bottom nine deciles of baseline risk beliefs; fatalistic people are those in the top decile.

† Differences and p-values in columns 3 and 7 are adjusted for sampling strata and clustered by village: * $p < 0.01$; ** $p < 0.05$; *** $p < 0.1$.

Table 2
Effect of Treatment on Self-Reported HIV Testing Decisions

	Outcome: Tested for HIV Since Baseline					
	(1)	(2)	(3)	(4)	(5)	(6)
(1) Fatalistic (10th Decile of Baseline Beliefs)	-0.003	-0.008	0.126*	-0.080*	-0.084*	-0.090**
	(0.038)	(0.039)	(0.071)	(0.043)	(0.042)	(0.041)
(2) Treatment		-0.030			-0.046*	-0.038
		(0.022)			(0.024)	(0.023)
(3) (Treatment) × (Fatalistic)					0.205**	0.282***
					(0.082)	(0.098)
Control-group Data	Yes	Yes	No	Yes	Yes	Yes
Treatment-group Data	Yes	Yes	Yes	No	Yes	Yes
T Interacted w/ Other Baseline Covariates	No	No	No	No	No	Yes
Observations	1,083	1,083	543	540	1,083	1,044
Adjusted R-squared	0.011	0.012	0.030	0.002	0.018	0.028
Control-group Mean	0.152	0.152	0.137	0.152	0.152	0.151
Control-group SD	0.359	0.359	0.344	0.359	0.359	0.359
Treatment Effect for Fatalistic People (2 + 3)					0.159	0.245
					(0.074)	(0.091)
Fatalistic vs. Non-Fatalistic Difference for Treatment Group (1 + 3)					0.121	0.192
					(0.070)	(0.089)

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. All regressions include controls for stratification cell fixed effects. Main effects are included for all variables included in interactions. Heteroskedasticity-robust standard errors, clustered by village, in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Figure 1
Estimated Treatment Effects by Decile of Baseline Risk Beliefs

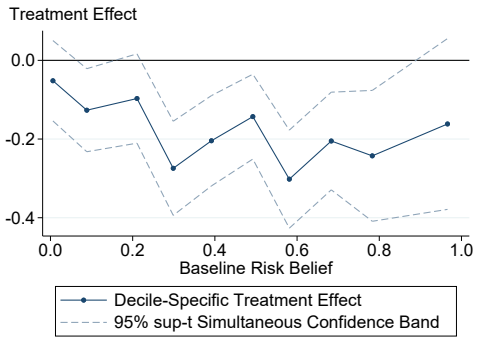

Panel A: Effects on Endline Risk Beliefs

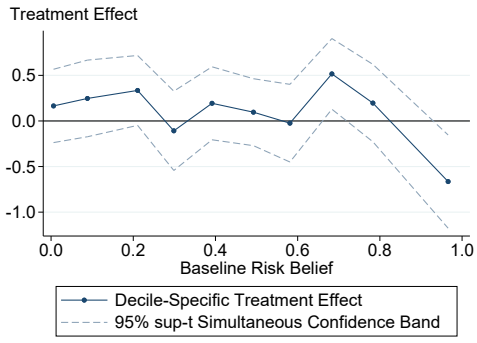

Panel B: Effects on Endline Sexual Activity

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Transmission Rate Belief) $\times$ (Prevalence Belief). The x-axis shows the mean value of baseline risk beliefs for each decile. Sexual Activity is measured as the log of sex acts in the past week, using the Ravallion (2017) transformation. The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. Confidence bands are clustered by village.

Figure 2
Nonparametric Estimates of Treatment Effects by Baseline Risk Belief

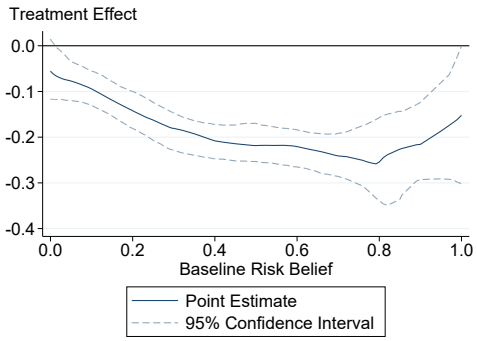

Panel A: Effects on Endline Risk Beliefs

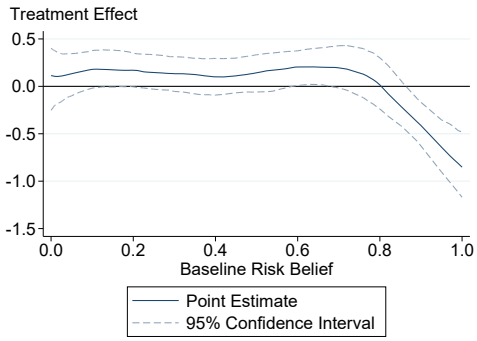

Panel B: Effects on Endline Sexual Activity

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Transmission Rate Belief) $\times$ (Prevalence Belief). Sexual Activity is measured as the log of sex acts in the past week, using the Ravallion (2017) transformation. The y-axis shows non-parametric estimates of the treatment effect for each value of Baseline Risk Beliefs using the local linear regression method described in Section 2.1. Confidence intervals are clustered by village.

Figure 3
Treatment Effect Heterogeneity by Deciles of Baseline Risk Beliefs
and Baseline Sexual Activity
Outcome: Ln(Sex Acts in Past 7 Days)

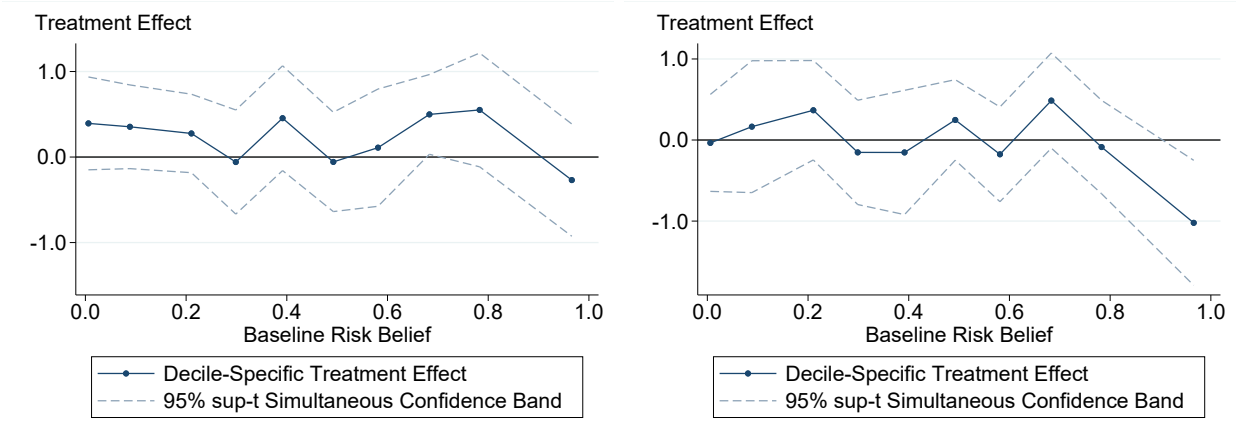

Panel A: Below-Median Sexual Activity at Baseline

Panel B: Above-Median Sexual Activity at Baseline

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.
The difference in the treatment effects for the top decile is 0.752 log points (SE = 0.368, $p = 0.045$).

Figure 4
Treatment Effect Heterogeneity by Deciles of Baseline Risk Beliefs
and Self-Reported HIV Status
Outcome: Ln(Sex Acts in Past 7 Days)

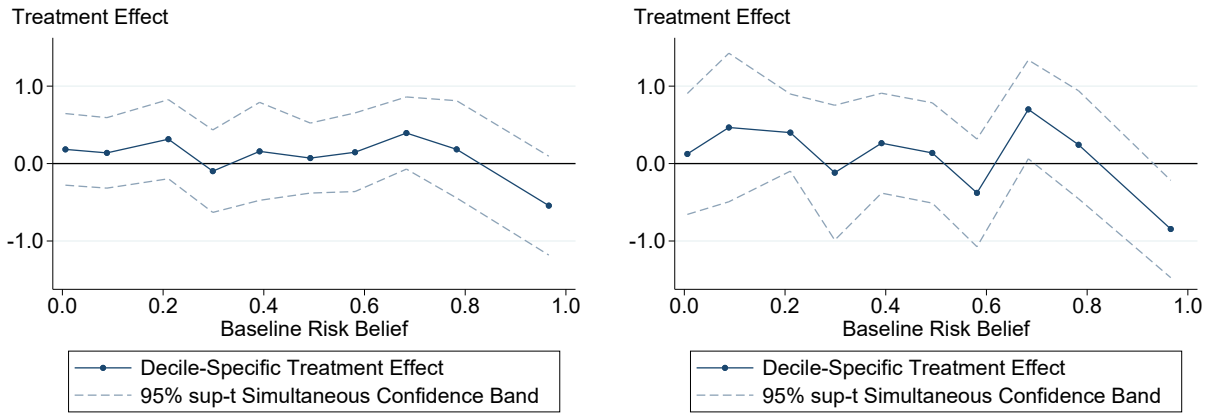

Panel A: No Chance I am HIV-positive

Panel B: Any Chance I am HIV-positive

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.
The difference in the treatment effects for the top decile is 0.302 log points (SE = 0.293,p = 0.307).

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.

Akesson, Jesper, Sam Ashworth-Hayes, Robert Hahn, Robert D. Metcalfe, and Itzhak Rasooly. 2022. "Fatalism, Beliefs, and Behaviors During the COVID-19 Pandemic." *Journal of Risk and Uncertainty*, forthcoming.

Altman, Douglas G., and J. Martin Bland. 1995. "Statistics Notes: Absence of Evidence Is Not Evidence of Absence." *BMJ*, 311(7003): 485.

Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber. 2005. "An evaluation of instrumental variable strategies for estimating the effects of catholic schooling." *Journal of Human Resources*, 40(4): 791–821. Publisher: University of Wisconsin Press.

Anderson, Michael L. 2008. "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association*, 103(484): 1481–1495.

Anglewicz, Philip, and Hans-Peter Kohler. 2009. "Overestimating HIV Infection: The Construction and Accuracy of Subjective Probabilities of HIV Infection in Rural Malawi." *Demographic Research*, 20: 65–96.

Baird, Sarah J., Richard S. Garfein, Craig T. McIntosh, and Berk Özler. 2012. "Effect of a Cash Transfer Programme for Schooling on Prevalence of HIV and Herpes Simplex Type 2 in Malawi: A Cluster Randomized Trial." *The Lancet*, 1–10.

Baranov, Victoria, and Hans-Peter Kohler. 2018. "The Impact of AIDS Treatment on Savings and Human Capital Investment in Malawi." *American Economic Journal: Applied Economics*, 10(1): 266–306.

Barnett, Tony, and Piers Blaikie. 1992. *AIDS in Africa: Its Present and Future Impact.* London:The Guilford Press.

Belli, Robert F., William L. Shay, and Frank P. Stafford. 2001. "Event History Calendars and Question List Surveys: A Direct Comparison of Interviewing Methods." *Public Opinion Quarterly*, 65(1): 45–74.

Bennear, Lori, Alessandro Tarozzi, Alexander Pfaff, Soumya Balasubramanya, Kazi Matin Ahmed, and Alexander van Geen. 2013. "Impact of a Randomized Controlled Trial in Arsenic Risk Communication on Household Water-Source Choices in Bangladesh." *Journal of Environmental Economics and Management*, 65(2): 225–240.

Bolstad, William M, and James M Curran. 2017. *Introduction to Bayesian Statistics.* 3rd ed., Hoboken, New Jersey:John Wiley & Sons, Inc.

Boozer, M. A, and T. J Philipson. 2000. "The Impact of Public Testing for Human Immunodeficiency Virus." *Journal of Human Resources*, 419–446.

Bruhn, Miriam, and David McKenzie. 2009. "In Pursuit of Balance: Randomization in Practice in Development Field Experiments." *American Economic Journal: Applied Economics*, 1(4): 200–232.

Cohen, Myron S., Ying Q. Chen, Marybeth McCauley, Theresa Gamble, Mina C. Hosseinipour, Nagalingeswaran Kumarasamy, James G. Hakim, Johnstone Kumwenda, Beatriz Grinsztejn, Jose H.S. Pilotto, Sheela V. Godbole, Sanjay Mehendale, Suwat Chariyalertsak, Breno R. Santos, Kenneth H. Mayer, Irving F. Hoffman, Susan H. Eshleman, Estelle Piwowar-Manning, Lei Wang, Joseph Makhema, Lisa A. Mills, Guy de Bruyn, Ian Sanne, Joseph Eron, Joel Gallant, Diane Havlir, Susan Swindells, Heather Ribaudo, Vanessa Elharrar, David Burns, Taha E. Taha, Karin Nielsen-Saines, David Celentano, Max Essex, and Thomas R. Fleming. 2011. "Prevention of HIV-1 Infection with Early Antiretroviral Therapy." *New England Journal of Medicine*, 365(6): 493–505.

Conroy, Amy A. 2014. "Marital Infidelity and Intimate Partner Violence in Rural Malawi: A Dyadic Investigation." *Archives of Sexual Behavior*, 43(7): 1303–1314.

Corno, Lucia, and Áureo De Paula. 2015. "Risky sexual behavior: biological markers and self-reported data." Working Paper.

Cowen, Tyler. 2020. "Being Irresponsible About Covid Isn't the Main Problem." *Bloomberg Opinion*.

Delavande, Adeline, and Hans-Peter Kohler. 2012. "The Impact of HIV Testing on Subjective Expectations and Risky Behavior in Malawi." *Demography*, 49(3): 1011–1036.

Delavande, Adeline, Xavier Giné, and David McKenzie. 2011. "Measuring Subjective Expectations in Developing Countries: A Critical Review and New Evidence." *Journal of Development Economics*, 94(2): 151–163.

Derksen, Laura, Adamson Muula, and Joep van Oosterhout. 2022. "Love in the Time of HIV: Testing as a Signal of Risk." Working Paper.

Dow, William H., Tomas J. Philipson, and Xavier Sala-i Martin. 1999. "Longevity Complementarities under Competing Risks." *American Economic Review*, 89(5): 1358–1371.

Dupas, Pascaline. 2011. "Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya." *American Economic Journal: Applied Economics*, 3(1): 1–34.

Ferrer, Rebecca A, and William MP Klein. 2015. "Risk Perceptions and Health Behavior." *Current Opinion in Psychology*, 5: 85–89.

Frison, Lars, and Stuart J. Pocock. 1992. "Repeated Measures in Clinical Trials: Analysis Using Mean Summary Statistics and Its Implications for Design." *Statistics in Medicine*, 11(13): 1685–1704.

Gayer, T., J.T. Hamilton, and W.K. Viscusi. 2002. "The Market Value of Reducing Cancer Risk: Hedonic Housing Prices with Changing Information." *Southern Economic Journal*, 266–289.

Gerrard, Meg, Frederick X. Gibbons, Alida C. Benthin, and Robert M. Hessling. 1996. "A Longitudinal Study of the Reciprocal Nature of Risk Behaviors and Cognitions in Adolescents: What You Do Shapes What You Think, and Vice Versa." *Health Psychology*, 15(5): 344. Publisher: US: American Psychological Association.

Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin Expected Utility with Non-Unique Prior." *Journal of Mathematical Economics*, 18(2): 141–153.

Godlonton, Susan, Alister Munthali, and Rebecca Thornton. 2016. "Responding to Risk: Circumcision, Information, and HIV Prevention." *The Review of Economics and Statistics*, 98(2): 333–349.

Gong, Erick. 2015. "HIV Testing & Risky Sexual Behaviour." *Economic Journal*, 125(582): 32–60.

Goyal, Stuti, and Jack Cavanagh. 2024. "Celebrating a Decade of the AEA RCT Registry."

Gray, Ronald H, Godfrey Kigozi, David Serwadda, Frederick Makumbi, Stephen Watya, Fred Nalugoda, Noah Kiwanuka, Lawrence H Moulton, Mohammad A Chaudhary, Michael Z Chen, Nelson K Sewankambo, Fred Wabwire-Mangen, Melanie C Bacon, Carolyn F M Williams, Pius Opendi, Steven J Reynolds, Oliver Laeyendecker, Thomas C Quinn, and Maria J Wawer. 2007. "Male circumcision for HIV prevention in men in Rakai, Uganda: a randomised trial." *The Lancet*, 369: 657–666.

Greenwood, Jeremy, Philipp Kircher, Cezar Santos, and Michèle Tertilt. 2019. "An Equilibrium Model of the African HIV/AIDS Epidemic." *Econometrica*, 87(4): 1081–1113.

Hollingsworth, T. Déirdre, Roy M. Anderson, and Christophe Fraser. 2008. "HIV-1 Transmission, by Stage of Infection." *Journal of Infectious Diseases*, 198(5): 687–693.

Hudomiet, Peter, Gábor Kézdi, and Robert J. Willis. 2011. "Stock Market Crash and Expectations of American Households." *Journal of Applied Econometrics*, 26(3): 393–415.

Imbens, Guido W., and Donald B. Rubin. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences.* Cambridge University Press.

Johnson, Weldon T., and John D. Delamater. 1976. "Response Effects in Sex Surveys." *Public Opinion Quarterly*, 40(2): 165–181.

Kadyoma, Frank, Mumbi Mbugua, Agnes Ndung'u, and Edward Omae. 2012. *Malawi Junior Secondary Life Skills: Form 2.* Nairobi, Kenya:Longhorn Publishers.

Kaler, Amy. 2003. "'My Girlfriends Could Fill a Yanu-Yanu Bus': Rural Malawian Men's Claims About Their Own Serostatus." *Demographic Research*, Special Collection(1): 349–372.

Kaler, Amy, and Susan Watkins. 2010. "Asking God About the Date You Will Die: HIV Testing as a Zone of Uncertainty in Rural Malawi." *Demographic Research*, 23: 905–932.

Kerwin, Jason T. 2012. "'Rational Fatalism': Non-Monotonic Choices in Response to Risk." Working Paper, WGAPE Spring 2012. Available at https://cega.berkeley.edu/assets/cega_events/4/Kerwin-Jason_Rational-Fatalism_2012.pdf.

Kerwin, Jason T., and Natalia Ordaz Reynoso. 2021. "You Know What I Know: Interviewer Knowledge Effects in Subjective Expectation Elicitation." *Demography*, 58(1): 1–29.

Koopman, James S., Carl P. Simon, and Chris P. Riolo. 2005. "When to Control Endemic Infections by Focusing on High-Risk Groups." *Epidemiology*, 16(5): 621–627.

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

Luke, Nancy, Shelley Clark, and Eliya M. Zulu. 2011. "The Relationship History Calendar: Improving the Scope and Quality of Data on Youth Sexual Behavior." *Demography*, 48(3): 1151–1176.

Malawi National AIDS Commission. 2003. "National HIV/AIDS Policy: A Call for Renewed Action." Malawi National AIDS Commission (NAC), Lilongwe.

Malawi National AIDS Commission. 2009. "National HIV/AIDS Prevention Strategy: 2009 to 2013." Malawi National AIDS Commission (NAC), Lilongwe.

Malawi National Statistical Office, and ORC-MACRO. 2010. "Malawi Demographic and Health Survey 2010." Calverton, MD: MEASURE DHS.

Manski, Charles F. 2004. "Measuring Expectations." *Econometrica*, 72(5): 1329–1376.

Matthies, Konstantin, and Flavio Toxvaerd. 2023. "Rather Doomed Than Uncertain: Risk Attitudes and Transmissive Behavior Under Asymptomatic Infection." *Economic Theory*, 76(1): 1–44.

Montiel Olea, José Luis, and Mikkel Plagborg-Møller. 2019. "Simultaneous Confidence Bands: Theory, Implementation, and an Application to SVARs." *Journal of Applied Econometrics*, 34(1): 1–17. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jae.2656.

O'Donoghue, Ted, and Matthew Rabin. 2001. "Risky Behavior Among Youths: Some Issues from Behavioral Economics." In *Risky Behavior among Youths: An Economic Analysis*. , ed. Jonathan Gruber, 29–68. Chicago, IL:University of Chicago Press.

Office of the Secretary. 1979. "The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research." National Commission for the Protection of Human Subjects of Biomedical Behavioral Research, Bethesda, MD.

Oster, Emily. 2012. "HIV and Sexual Behavior Change: Why Not Africa?" *Journal of Health Economics*, 31(1): 35–49.

Oster, Emily. 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business & Economic Statistics*, 37(2): 187–204. Publisher: Taylor & Francis.

Oster, Emily. 2020. "The 'Just Stay Home' Message Will Backfire." *The Atlantic*.

Peltzman, Sam. 1975. "The Effects of Automobile Safety Regulation." *Journal of Political Economy*, 83(4): 677–725.

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

Ravallion, Martin. 2017. "A Concave Log-Like Transformation Allowing Non-Positive Values." *Economics Letters*, 161: 130–132.

Rubin, Donald B. 1981. "The Bayesian Bootstrap." *The Annals of Statistics*, 9(1,): 130–134.

Schatz, Enid. 2005. "'Take Your Mat and Go!': Rural Malawian Women's Strategies in the HIV/AIDS Era." *Culture, Health & Sexuality*, 7(5): 479–492.

Shapiro, Joel, and Stephen Wu. 2011. "Fatalism and Savings." *The Journal of Socio-Economics*, 40(5): 645–651.

Shrestha, Maheshwor. 2020. "Get Rich or Die Tryin': Perceived Earnings, Perceived Mortality Rates, and Migration Decisions of Potential Work Migrants from Nepal." *World Bank Economic Review*, 34(1): 1–27.

Simmons, Joseph P., Leif D. Nelson, and Uri Simonsohn. 2011. "False-Positive Psychology Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant." *Psychological Science*, 22(11): 1359–1366.

Swidler, Ann, and Susan Cotts Watkins. 2007. "Ties of Dependence: AIDS and Transactional Sex in Rural Malawi." *Studies in Family Planning*, 38(3): 147–162.

Tawfik, Linda, and Susan Cotts Watkins. 2007. "Sex in Geneva, Sex in Lilongwe, and Sex in Balaka." *Social Science & Medicine (1982)*, 64(5): 1090–1101.

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

Thornton, Rebecca L. 2012. "HIV Testing, Subjective Beliefs and Economic Behavior." *Journal of Development Economics*, 99(2): 300–313.

Viscusi, W. Kip. 1990. "Do Smokers Underestimate Risks?" *Journal of Political Economy*, 98(6): 1253–1269.

Wawer, Maria J., Ronald H. Gray, Nelson K. Sewankambo, David Serwadda, Xianbin Li, Oliver Laeyendecker, Noah Kiwanuka, Godfrey Kigozi, Mohammed Kiddugavu, Thomas Lutalo, Fred Nalugoda, Fred Wabwire-Mangen, Mary P. Meehan, and Thomas C. Quinn. 2005. "Rates of HIV-1 Transmission Per Coital Act, by Stage of HIV-1 Infection, in Rakai, Uganda." *Journal of Infectious Diseases*, 191: 1403–1409.

Weinstein, Neil D. 1989. "Optimistic Biases About Personal Risks." *Science*, 246(4935): 1232–1233. Publisher: American Association for the Advancement of Science.

Wilson, Nicholas L., Wentao Xiong, and Christine L. Mattson. 2014. "Is Sex Like Driving? HIV Prevention and Risk Compensation." *Journal of Development Economics*, 106: 78–91.

Supplemental Online Appendix, Not Intended for Publication

A1 Appendix Tables and Figures

Appendix Table A1
Baseline Balance for the Full Sample

	Panel A: Control (C) vs. Treatment (T)				Panel B: Non-Fatalistic (N) vs. Fatalistic (F)			
	C Mean (SD)	T Mean (SD)	Diff.† (p-value)	Obs.	N Mean (SD)	F Mean (SD)	Diff.† (p-value)	Obs.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sexual Activity								
Any Sex in Past Week	0.541	0.507	-0.036	1,292	0.522	0.571	0.046	1,275
	(0.499)	(0.500)	(0.111)		(0.500)	(0.497)	(0.309)	
Total Acts in Past Week	1.798	1.615	-0.185	1,292	1.705	1.839	0.113	1,275
	(2.471)	(2.380)	(0.155)		(2.440)	(2.403)	(0.618)	
Unprotected Acts in Past Week	1.569	1.471	-0.100	1,292	1.509	1.723	0.190	1,275
	(2.376)	(2.323)	(0.446)		(2.353)	(2.413)	(0.399)	
Sex Partners in Past 30 Days	0.818	0.797	-0.024	1,290	0.804	0.856	0.048	1,273
	(0.498)	(0.762)	(0.515)		(0.656)	(0.537)	(0.316)	
Condoms Acquired in Past 30 Days	4.739	3.530	-1.205	1,288	3.989	6.188	2.200	1,271
	(15.003)	(11.549)	(0.122)		(12.921)	(18.218)	(0.314)	
Years Sexually Active	13.100	13.204	0.117	1,275	12.940	15.330	2.445**	1,261
	(8.279)	(8.603)	(0.815)		(8.366)	(8.923)	(0.021)	
Lifetime Sex Partners	3.117	3.557	0.414**	1,288	3.289	3.811	0.442	1,272
	(2.684)	(4.734)	(0.042)		(3.861)	(3.798)	(0.203)	
Any Chance of Having HIV	0.344	0.352	0.008	1,277	0.343	0.387	0.051	1,261
	(0.475)	(0.478)	(0.788)		(0.475)	(0.489)	(0.305)	
Overall Sexual Activity Index	0.028	-0.028	-0.059	1,277	-0.005	0.091	0.089	1,261
	(0.997)	(1.003)	(0.266)		(1.004)	(0.971)	(0.344)	
Demographics								
Male	0.425	0.436	0.000	1,292	0.431	0.455	-0.000	1,275
	(0.495)	(0.496)	(1.000)		(0.495)	(0.500)	(1.000)	
Married	0.829	0.803	-0.025	1,290	0.817	0.839	0.023	1,273
	(0.377)	(0.398)	(0.316)		(0.387)	(0.369)	(0.575)	
Age	29.133	29.589	0.465	1,292	29.101	31.804	2.789***	1,275
	(8.417)	(8.333)	(0.339)		(8.296)	(8.836)	(0.009)	
Years of Education	5.758	5.858	0.097	1,292	5.794	6.000	0.190	1,275
	(3.347)	(3.484)	(0.723)		(3.397)	(3.604)	(0.600)	
Household Size	5.039	4.870	-0.176	1,292	4.943	5.089	0.159	1,275
	(2.237)	(2.036)	(0.254)		(2.122)	(2.331)	(0.439)	
Spending in Past 30 Days	292.390	293.010	1.698	1,292	290.939	310.361	20.771	1,275
	(383.593)	(572.544)	(0.954)		(486.888)	(512.716)	(0.696)	
Assets Owned	4.188	3.937	-0.248	1,292	4.043	4.259	0.204	1,275
	(2.427)	(2.311)	(0.192)		(2.348)	(2.564)	(0.382)	
Ravens Score [0-3]	1.551	1.538	-0.019	1,291	1.543	1.563	-0.004	1,275
	(0.989)	(1.002)	(0.766)		(0.998)	(0.994)	(0.962)	
Numeracy [0-3]	0.715	0.818	0.096*	1,292	0.774	0.732	-0.057	1,275
	(0.929)	(1.007)	(0.095)		(0.967)	(1.004)	(0.529)	
Chance of Winning Question	0.219	0.249	0.027	1,292	0.234	0.241	0.003	1,275
	(0.414)	(0.433)	(0.283)		(0.423)	(0.430)	(0.944)	
Risk Attitude	0.261	0.274	0.014	1,288	0.267	0.239	-0.032	1,271
	(0.440)	(0.447)	(0.634)		(0.442)	(0.428)	(0.440)	
Christian	0.910	0.927	0.017	1,292	0.917	0.955	0.042**	1,275
	(0.286)	(0.260)	(0.472)		(0.277)	(0.207)	(0.046)	
Muslim	0.085	0.060	-0.025	1,292	0.074	0.045	-0.031	1,275
	(0.280)	(0.238)	(0.281)		(0.262)	(0.207)	(0.130)	

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. In Panel B, fatalistic people are those in the top decile of baseline risk beliefs.
† Differences and p-values in columns 3 and 7 are adjusted for sampling strata and clustered by village: * $p < 0.01$; ** $p < 0.05$; *** $p < 0.1$.

Appendix Table A2

Attrition Patterns by Sexual Activity

	Present in Final Sample[†]											
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treatment Group [T]	0.020	0.020	0.020	0.020	0.020	0.020	0.021	0.020	0.022	0.015	0.020	0.019
	(0.019)	(0.019)	(0.019)	(0.019)	(0.020)	(0.020)	(0.020)	(0.019)	(0.020)	(0.019)	(0.020)	(0.020)
T × (Any Sex in Past Week)		-0.001										0.154
		(0.034)										(0.109)
T × (Sex Acts in Past Week)			-0.013*									-0.038*
			(0.008)									(0.023)
T × (Unprotected Acts in Past Week)				-0.012								0.020
				(0.007)								(0.028)
T × (Sex Acts in Past 30 Days)					0.001							0.006
					(0.001)							(0.004)
T × (Sex Partners in Past 30 Days)						0.007						0.039
						(0.036)						(0.060)
T × (Condoms Acquired in Past 30							0.000					0.000
							(0.001)					(0.001)
T × (Years Sexually Active)								0.001				-0.001
								(0.002)				(0.005)
T × (Lifetime Sex Partners)									0.011*			0.011*
									(0.006)			(0.006)
T × (Any Chance of Having HIV)										0.017		0.010
										(0.031)		(0.034)
T × (Overall Sexual Activity Index)											-0.007	-0.080
											(0.020)	(0.107)
Observations	1,503	1,503	1,503	1,503	1,487	1,497	1,495	1,486	1,495	1,481	1,483	1,447
Adjusted R-squared	0.030	0.029	0.031	0.030	0.029	0.029	0.030	0.028	0.034	0.027	0.029	0.030
Control-group Mean	0.850	0.850	0.850	0.850	0.852	0.852	0.852	0.848	0.851	0.855	0.851	0.853

Notes: Present in Final Sample denotes the set of respondents who were contacted at baseline, had a complete baseline survey, and were subsequently found for the endline survey. Sample includes 1,503 sexually active adults who were successfully interviewed at baseline; 56 of these have missing data for at least one of the controls. All covariates are de-meaned prior to running the regression. Whenever regressions include an interaction between a covariate and the treatment, a main effect is included as well. Results in column 12 are from a regression that includes all interaction terms and main effects from column 13 of Appendix Table A3 as well. Heteroskedasticity-robust standard errors, clustered by village, in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Appendix Table A3
Attrition Patterns by Demographics

	Present in Final Sample[†]												
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Treatment Group [T]	0.020	0.020	0.020	0.019	0.020	0.020	0.020	0.020	0.021	0.021	0.020	0.019	0.019
	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.020)	(0.020)	(0.020)	(0.020)	(0.019)	(0.020)
T × (Male)	0.031												-0.013
	(0.034)												(0.045)
T × (Married)		-0.076											-0.110*
		(0.051)											(0.061)
T × (Age)			0.001										0.002
			(0.002)										(0.005)
T × (Years of education)				0.004									-0.002
				(0.006)									(0.007)
T × (Number of people in HH)					-0.002								-0.001
					(0.008)								(0.008)
T × (Spending in Past 30 Days)						0.000							0.000
						(0.000)							(0.000)
T × (# Assets Owned)							0.007						0.006
							(0.008)						(0.009)
T × (Ravens Score [0-3])								0.008					-0.003
								(0.019)					(0.021)
T × (Numeracy [0-3])									0.025				0.006
									(0.021)				(0.031)
T × (Chance of Winning Question)										0.062			0.029
										(0.045)			(0.064)
T × (Risk Attitude)											0.003		0.008
											(0.042)		(0.044)
T × (Christian)												-0.001	-0.032
												(0.084)	(0.083)
Observations	1,503	1,501	1,503	1,503	1,503	1,503	1,503	1,498	1,499	1,499	1,495	1,503	1,447
Adjusted R-squared	0.030	0.032	0.029	0.029	0.029	0.029	0.029	0.028	0.030	0.034	0.028	0.030	0.030
Control-group Mean	0.850	0.850	0.850	0.850	0.850	0.850	0.850	0.852	0.852	0.852	0.852	0.850	0.853

Notes: Present in Final Sample denotes the set of respondents who were contacted at baseline, had a complete baseline survey, and were subsequently found for the endline survey. Sample includes 1,503 sexually active adults who were successfully interviewed at baseline; 56 of these have missing data for at least one of the controls. All covariates are de-meaned prior to running the regression. Whenever regressions include an interaction between a covariate and the treatment, a main effect is included as well. Muslim is omitted from the set of demographic controls due to collinearity with Christian. Results in column 13 are from a regression that includes all interaction terms and main effects from column 12 of Appendix Table A2 as well. Heteroskedasticity-robust standard errors, clustered by village, in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Appendix Table A4

Average Treatment Effects on HIV Risk Beliefs

	Perceived HIV Transmission Rate, if Partner Infected				Perceived HIV Prevalence		Composite Beliefs: P(Contract HIV from Unpro. Sex w/Random Attractive Person‡)		Any Chance of Having HIV	Any Chance of Partner Having HIV
	One Act		One Year†		All Local People‡	Attractive Local People‡	One Act	One Year†		
	Unprotected	w/Condom	Unprotected	w/Condom						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment Group	-0.384***	-0.045***	-0.371***	-0.071***	-0.162***	-0.047***	-0.182***	-0.185***	-0.021	-0.012
	(0.019)	(0.006)	(0.016)	(0.012)	(0.016)	(0.015)	(0.014)	(0.015)	(0.029)	(0.028)
Observations	1,281	1,283	1,276	1,276	1,257	1,254	1,252	1,251	1,242	1,229
Adjusted R-squared	0.315	0.066	0.328	0.142	0.157	0.081	0.200	0.182	0.184	0.230
Control-group Mean	0.743	0.0819	0.906	0.177	0.487	0.464	0.351	0.424	0.362	0.368
Control-group SD	0.317	0.162	0.196	0.264	0.289	0.265	0.268	0.263	0.481	0.483

Notes: † For couples having typical sexual behavior over the course of one year.
‡ Prevalence belief variables are questions specifically about members of the opposite sex.
Treatment effects estimated by regressing endline beliefs on the treatment indicator, controlling for baseline beliefs and stratification cell indicators: $x_i^e = \beta_0 + \beta_1 T_i + \beta_2 x_i^b + Z_i'\eta + \varepsilon_i$. Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Heteroskedasticity-robust standard errors, clustered by village, in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Appendix Table A5

Heterogeneity in Treatment Effects by Deciles of Baseline Risk Beliefs

	Outcome: Log Sex Acts in Past Week		
	(1)	(2)	(3)
Treatment [T] X			
1st Decile of Baseline Risk Beliefs	0.165	0.045	0.086
	(0.144)	(0.185)	(0.181)
2nd Decile of Baseline Risk Beliefs	0.247	0.120	0.111
	(0.150)	(0.194)	(0.191)
3rd Decile of Baseline Risk Beliefs	0.335**	0.360**	0.333**
	(0.137)	(0.154)	(0.154)
4th Decile of Baseline Risk Beliefs	-0.108	0.017	-0.006
	(0.156)	(0.180)	(0.179)
5th Decile of Baseline Risk Beliefs	0.194	0.109	0.115
	(0.143)	(0.170)	(0.175)
6th Decile of Baseline Risk Beliefs	0.096	0.170	0.195
	(0.131)	(0.145)	(0.153)
7th Decile of Baseline Risk Beliefs	-0.024	0.053	0.042
	(0.152)	(0.166)	(0.170)
8th Decile of Baseline Risk Beliefs	0.516***	0.460***	0.432**
	(0.139)	(0.168)	(0.176)
9th Decile of Baseline Risk Beliefs	0.196	0.022	0.050
	(0.152)	(0.170)	(0.166)
10th Decile of Baseline Risk Beliefs	-0.665***	-0.756***	-0.743***
	(0.183)	(0.187)	(0.185)
Control for BL Outcome	Yes	No	No
Stratification Cell FEs	Yes	Yes	No
T Interacted w/BL Outcome	Yes	No	No
T Interacted with Other Baseline Covariates	Yes	No	No
Observations	1,232	1,275	1,275
Adjusted R-squared	0.307	0.031	0.016
Control-group Mean	1.683	1.673	1.673
Control-group SD	2.390	2.382	2.382

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). Heteroskedasticity-robust standard errors, clustered by village, in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Appendix Table A6

Means and SDs of Sexual Activity by Study Arm, Survey Wave, and Decile of Baseline Beliefs

	Baseline			Endline			Change from Baseline to Endline			
	C Mean (SD)	T Mean (SD)	Diff. (p-value)	C Mean (SD)	T Mean (SD)	Diff. (p-value)	C Mean (SD)	T Mean (SD)	Diff. (p-value)	N
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Decile of Baseline Risk Beliefs										
1st	1.855	1.237	-0.617	1.364	1.562	0.199	-0.491	0.325	0.816**	135
	(2.870)	(2.094)	(0.174)	(1.985)	(2.134)	(0.605)	(2.892)	(1.960)	(0.040)	
2nd	1.907	1.152	-0.755	1.326	1.481	0.155	-0.581	0.329	0.911**	122
	(3.123)	(1.895)	(0.208)	(2.589)	(2.353)	(0.765)	(1.679)	(2.520)	(0.017)	
3rd	1.692	1.844	0.152	1.308	2.221	0.913**	-0.385	0.377	0.761**	142
	(2.229)	(2.870)	(0.708)	(1.828)	(2.718)	(0.018)	(1.791)	(2.739)	(0.039)	
4th	1.403	1.792	0.389	1.419	1.434	0.015	0.016	-0.358	-0.375	115
	(1.954)	(2.713)	(0.378)	(2.021)	(1.995)	(0.968)	(2.053)	(2.725)	(0.378)	
5th	1.974	1.790	-0.183	1.724	2.048	0.325	-0.250	0.258	0.508	138
	(2.713)	(2.841)	(0.718)	(2.549)	(3.000)	(0.482)	(2.862)	(2.395)	(0.234)	
6th	1.573	1.549	-0.024	1.607	1.934	0.327	0.034	0.385	0.351	180
	(2.567)	(1.827)	(0.943)	(2.716)	(2.744)	(0.430)	(3.256)	(2.788)	(0.402)	
7th	1.698	1.982	0.283	2.095	2.091	-0.004	0.397	0.109	-0.288	118
	(2.061)	(2.513)	(0.444)	(2.716)	(2.351)	(0.992)	(1.931)	(2.132)	(0.441)	
8th	1.955	1.939	-0.016	1.388	2.306	0.918**	-0.567	0.367	0.935**	116
	(2.312)	(2.688)	(0.970)	(2.132)	(2.592)	(0.042)	(2.530)	(2.079)	(0.021)	
9th	2.333	1.565	-0.768	1.725	1.804	0.079	-0.608	0.239	0.847*	97
	(2.673)	(2.083)	(0.104)	(2.281)	(2.400)	(0.849)	(2.040)	(2.469)	(0.070)	
10th	1.913	1.721	-0.192	2.594	0.930	-1.664***	0.681	-0.791	-1.472***	112
	(2.331)	(2.539)	(0.686)	(2.475)	(1.502)	(0.000)	(2.552)	(2.077)	(0.002)	

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) × (Baseline Prevalence Belief). Table shows means and differences of the number of sex acts in the past week for each study arm and decile of baseline risk beliefs, with standard errors and cluster-adjusted p-values in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Appendix Figure A1

Perceived HIV Risk Survey Questions (Female Versions)

Belief Variable	**Question Text**
Perceived HIV Transmission Rate	
One Year, Unprotected	If 100 women, who do **not** have HIV, each have an HIV-positive sex partner for **one year,** and do **not** use condoms when having sex, how many of the women do you think will have HIV at the end of the year?
One Year, W/Condom	If 100 women, who do **not** have HIV, each have an HIV-positive sex partner for **one year,** and **do** use condoms when having sex, how many of the women do you think will have HIV at the end of the year?
One Act, Unprotected	If 100 women, who do **not** have HIV, each sleep with a man who is HIV positive
One Act, W/Condom	If 100 women, who do **not** have HIV, each sleep with a man who is HIV positive tonight and **do** use a condom, how many of them do you think will have HIV after the night?
Perceived HIV Prevalence	
All Local People	If we took a group of 100 men from this area - just normal men who you found working nearby or in homes - how many of them do you think would have HIV?
Attractive Local People	Think of ten men from your village who you think are attractive. How many of them do you think would have HIV?

Notes: Survey questions were gender-specific, so men were asked about 100 men and women were asked about 100 women. All survey questions were asked in Chichewa (translated versions available upon request).

Appendix Figure A2
Distributions of Baseline HIV Infection Risk Beliefs

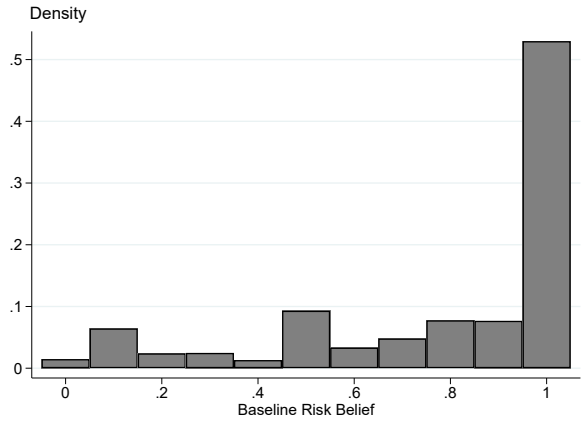

Panel A: Per-Act Infection Risk from Unprotected Sex with an Infected Partner

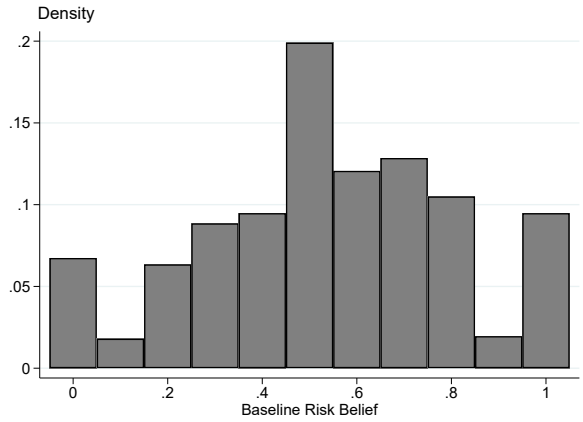

Panel B: Prevalence of HIV Among Attractive Local People

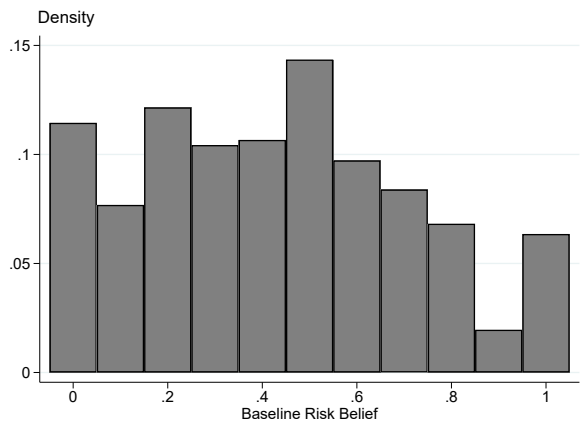

Panel C: Per-Act Infection Risk from Unprotected Sex with a Randomly Selected Partner

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed.

Appendix Figure A3
Distributions of Endline HIV Infection Risk Beliefs by Study Arm

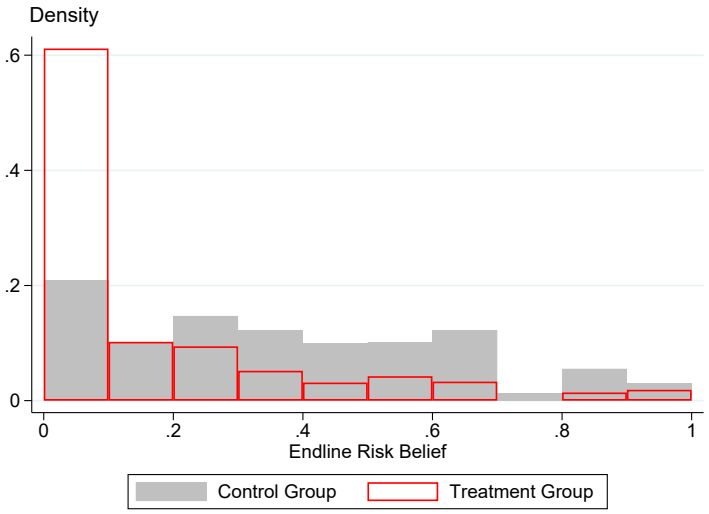

Per-Act Infection Risk from Unprotected Sex with a Randomly Selected Partner

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed.

Appendix Figure A4
Distributions of Endline Sexual Activity by Study Arm

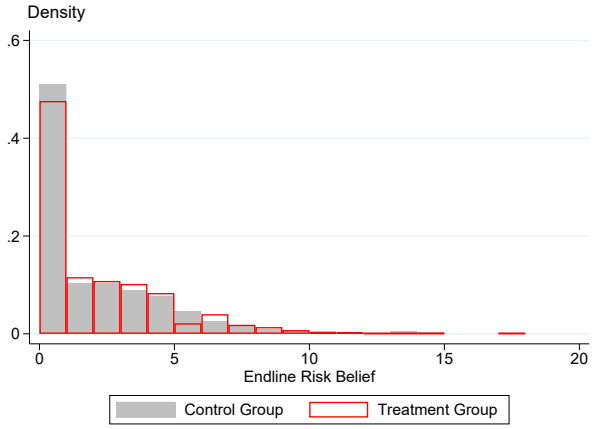

Panel A: Overall

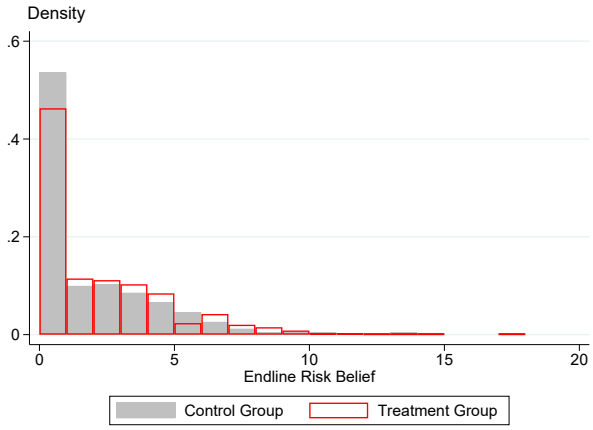

Panel B: Non-Fatalistic Sample

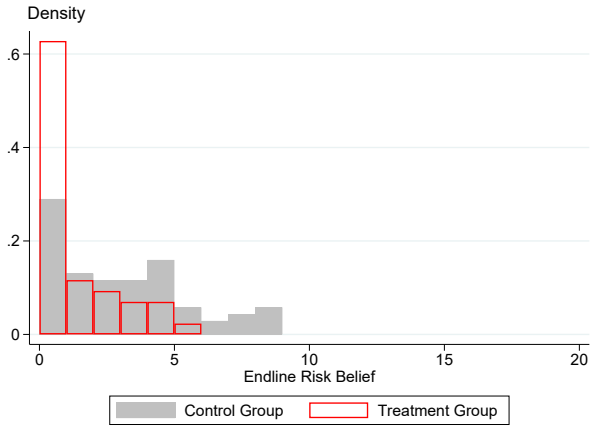

Panel C: Fatalistic Sample

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Overall sample (Panel A) includes every respondent. Non-fatalistic people (Panel B) are those in the bottom nine deciles of baseline risk beliefs; fatalistic people (Panel C) are those in the top decile.

Appendix Figure A5

Treatment Effects on Perceived Lifespan Conditional on Contracting HIV by Decile of Baseline Risk Beliefs

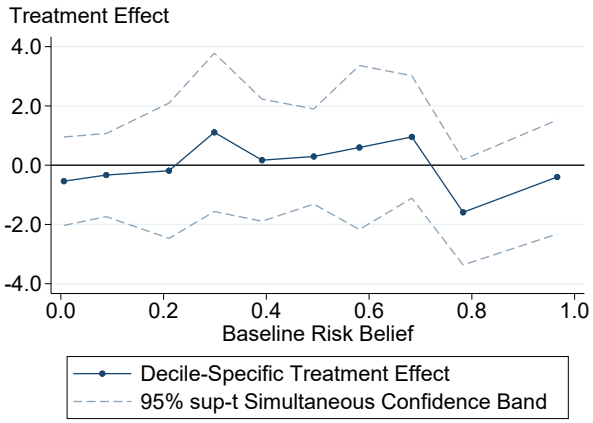

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.

B Details of Sampling Strategy

The village sample for the study was constructed from the Malawi National Statistics Office GIS files for the 2008 Census. I began by removing all duplicate village entries from the dataset.[20] Because existing evidence indicates that fatalistic responses to HIV risks and risky sexual activity may be concentrated around major trading centers (Kaler 2003), I then constructed sampling strata based on the distance to the closest major trading center.[21] 24 of the sampled villages (34%) were within 2 km of a trading center, another 24 (34%) were within 2 and 5 km from a trading center, and 22 (31%) were more than 5 km away from the closest center; this compares with overall proportions of 10%, 40% and 50% of all villages in TA Mwambo. In discussions with people from the local area, 2 km was generally agreed to be the maximal distance people will walk for nightlife. These strata thus roughly proxy for how easily people could access the trading centers in order to drink and search for sex partners. Within each sampling stratum, I randomly assigned half of the villages to the treatment group and half to the control group. Appendix Table B1 shows the distribution of respondents in each sampling stratum and study arm.

In each village, a team of enumerators first conducted a comprehensive household census. Using this census, 15 men and 15 women aged 18-49 were then sampled from each village, with only one respondent allowed per household. The sample was thus stratified by both gender and distance to the nearest trading center, so the effective sampling strata are formed by combinations of gender and distance indicators. Some villages had too few households for 15 age-eligible adults of each gender to be selected, and hence the maximum feasible number was chosen instead. The initial sample comprised 2,024 individuals. The survey team attempted to contact all sampled people for a baseline survey. Although refusals were rare ($< 1\%$ of respondents refused the baseline survey), 23% of sampled people could not be found at baseline, typically because they were temporarily away from the household; it

[20] The Population and Housing Census uses enumeration areas as its basic sampling unit, rather than villages. The boundaries of these enumeration areas commonly cross through villages, leading to duplicate entries in the GIS datasets.

[21] Trading centers were identified based on their designation by the 2008 Malawi Population and Housing Census. Since the study region (TA Mwambo) adjoins the city of Zomba, I also included the main markets in that city as trading center equivalents. In addition, based on conversations with knowledgeable locals, I included several more trading centers in the local area that were not designated as such by the census.

is common for people in this area of Malawi to travel during the agricultural off-season to look for casual wage labor. A total of 1,543 respondents had a successful baseline survey. Because the survey contained sensitive questions about sexual behavior, and the prediction of fatalistic responses holds solely for people who are sexually active, the survey used an early screening question to eliminate people who had never had sex from the sample. This removed 2.6% of the respondents, leaving 1,503 sexually active adults in the baseline survey.

Appendix Table B1
Sample Selection and Randomization

	Overall	Control	Treatment
Villages	70	35	35
Sampling Stratum[†]			
0-2 km from a trading center	24	12	12
2-5 km from a trading center	24	12	12
5+ km from a trading center	22	11	11
Respondents			
With Complete Baseline Survey	1503	759	744
With Complete Endline Survey	1292	645	647
Successful Followup Rate	0.86	0.85	0.87

C Ethical Considerations in Designing the Information Intervention

The key potential ethical concern about the design of this study was that on average people may react to HIV infection risks via conventional risk compensation. In this case the information treatment would increase the average amount of risky sex people have, leaving people in the treatment group worse off. This concern is mitigated by four factors. First, to the extent that we believe responsible adults can be trusted to make their own choices with the information they have, it is appropriate to provide people with better information rather than worse. The *de facto* policy in Malawi is to overstate HIV transmission risks. This strategy is potentially at odds with the first ethical principle emphasized in the Belmont Report, which is that individuals should be respected as autonomous persons:

> To respect autonomy is to give weight to autonomous persons' considered opinions and choices while refraining from obstructing their actions unless they are clearly detrimental to others. To show lack of respect for an autonomous agent is to repudiate that person's considered judgments, to deny an individual the freedom to act on those considered judgments, or to withhold information necessary to make a considered judgment, when there are no compelling reasons to do so.) (Office of the Secretary 1979)

Hence the policy of denying people information about the true risks they face is potentially unethical, given that there is very little empirical evidence that would provide compelling reasons to withhold that information.

Second, the information provided to the treatment group is medically accurate, publicly available information, drawn from research by Wawer et al. (2005). It is also the same information provided by the Malawi National AIDS Commission (NAC) in their policy documents. The *National HIV/AIDS Prevention Strategy: 2009 to 2013* states that the annual risk of HIV transmission for serodiscordant couples[22] is 5-10% (Malawi National AIDS Commission 2009, p.11); the figure I provide is at the upper end of this range, and so would be the least likely to induce conventional risk compensation. NAC's official policy is also that HIV information and education programs should provide accurate information about safer sex:

> Government, through the NAC, undertakes to do the following:
> • Ensure that all people have equal access to culturally sound and age-appropriate formal and nonformal HIV/AIDS information and education programmes, which shall include free and accurate information regarding mother-to-child transmission, breastfeeding, treatment, nutrition, change of lifestyle, safer sex and the importance of respect for and nondiscrimination against PLWAs [people living with AIDS].
> (Malawi National AIDS Commission 2003, [p.6])

Hence the additional information provided to the treatment group is completely consistent with Malawi government policy, and can be seen as a test of what would happen if HIV information and education campaigns actually provided HIV transmission risk information that is consistent with what NAC provides in reports that are available on its website.

[22] A couple where one partner is HIV-positive and the other is HIV-negative.

A third mitigating factor is that previous estimates of responses to HIV risks in Africa are very small in magnitude (e.g. Oster 2012), and the *ex ante* expected impact of the information treatment was small, limiting any potential harm. The reason that the experiment was still interesting was that the responses were not expected to be uniform. There is reason to believe that many people in Malawi may react fatalistically to HIV risks. Cross-sectional data from elsewhere in Zomba District shows suggestive evidence that the response of sexual behavior to HIV infection is positive for people with high risk beliefs (Kerwin 2012). Kaler (2003) documents that men from rural Southern Malawi employ fatalistic reasoning - saying that it is sometimes not worthwhile to use condoms, because the risk of contracting the virus is so high:

> And then I asked my in-law, "What do people do after noticing that his/her partner seems to have AIDS?" He said, "Some couples come to an end and for others the marriage continues." And I asked, "Do they use condoms then?" He said "I don't think they use [them] because it will just be a waste of time since both of them have contracted the disease." (Simon, journal May 3 2002)

For people who respond fatalistically, learning that their assessment of the risk is an overestimate will actually reduce sexual risk-taking, rather than increasing it. This experiment was designed to capture heterogeneity in responses around a mean response that is small in magnitude.

Finally, this concern is mitigated because excessively high risk beliefs are unlikely to persist in the long term. Serodiscordant couples are very common, and people can observe that it is possible for sexually active married couples to remain serodiscordant for a long time. This should cause them to update their risk beliefs downward, which would affect sexual behavior in a similar way to my information treatment, mitigating any net effects on sexual behavior in the long run. Moreover, if people realize that they were misled about the risks (or that their misconceptions were not corrected) they may lose trust in the medical and science community or the education system, and may also promulgate false rumors about HIV transmission and immunity. Since most people believe that the transmission rate of HIV is 100%, they may instead falsely assume that continued serodiscordance means that a specific person or group is immune to the virus. There is already evidence that the latter

is going on: 42% of my respondents said that they believed people with type-O blood were immune to HIV, an idea which has no basis in scientific fact.

A separate potential concern is that the information presented is about the approximate overall average risk, but transmission risks actually vary by demographic groups. For example, the transmission rate is 3 to 5 times higher for women than for men, and about 60% lower for circumcised men than for uncircumcised men. However, this concern is mitigated by the fact that baseline beliefs are very high (93% per year on average for the control group). Hence virtually all respondents in the treatment group have more-accurate beliefs after the information treatment than they did beforehand.

To ensure that respondents' well-being was protected, ethics oversight for this study was provided by both an in-country IRB (The University of Malawi College of Medicine Research and Ethics Committee, or COMREC) and one at my home institution (The University of Michigan's IRB-Health Sciences and Behavioral Sciences, or IRB-HSBS). The final study protocol (COMREC protocol # P.07/11/1107, IRB-HSBS protocol # HUM00052708) including the information treatment, was reviewed and approved by both IRBs. The approved protocol also included a management plan under which preliminary results were provided to the two IRBs in order to manage any possible rise in HIV transmissions as a result of the information treatment.

D Details of the Information Treatment

This section provides details of how the information treatment was presented to subjects in the study. The information treatment consisted of both an oral component and an interactive visual component.

The information treatment happened immediately after the baseline survey for treatment-group respondents. All participants were provided with basic information about the sexual transmission of HIV and the benefits of condoms. Knowledge of the basics of HIV transmission and prevention is already high in this population. In the 2010 DHS, nearly 100% of individuals said that HIV was sexually transmitted and over four fifths knew that condoms were effective prevention (Malawi National Statistical Office and ORC-MACRO 2010).

Treatment group respondents were also provided with information about the transmission rate of HIV, presented both orally and visually. I replicate the full information treatment, including the oral script and the visual diagrams, below as Appendix D.1.

In the oral component, the basic details of the original Rakai study were explained, with certain aspects simplified for clarity. Respondents were told that the study occurred in Uganda, and that 100 serodiscordant couples were followed for a single year.[23] They were told that all the couples had regular sex without using condoms, about once every three days on average, and asked how many people they thought would contract HIV. They were then informed that in fact only ten of the initially HIV-negative people became HIV-positive.[24] Respondents were asked if they believed the results of the study; enumerators were trained in how to respond to a number of common questions, such as whether the testing equipment was faulty.[25] The script listed the reasons that HIV transmission sometimes does not happen even when serodiscordant couples have unprotected sex, for example the fact that HIV sometimes cannot penetrate the genitalia. The script then emphasized that HIV transmission is something that happens by chance, comparing it to popular games of chance used by local cell phone companies as marketing tools.

The interactive visual component complemented the oral component and occurred at the same time. It involved showing respondents a diagram with 100 pairs of stick figures representing serodiscordant couples, with a black stick figure indicating an HIV-negative partner and white stick figure indicating an HIV-positive partner. The respondent was asked to guess the number of people who would contract HIV after a year of regular unprotected sex with an infected partner, and this guess was indicated by circling an appropriate number of these stick figure couples. When the true rate was presented, the enumerator showed a second diagram in which ten of the initially HIV-negative individuals had turned from black

[23] The Wawer et al. (2005) study includes 235 couples, 188 of which never used condoms when they had sex (results are not broken out by condom use, but condom use was very inconsistent and had no impact on the estimated transmission rate). Couples were observed over 10-month time windows, with some observed for multiple windows. I reduced this to 100 couples over the course of 1 year for clarity and simplicity.

[24] This is the annual transmission rate cited by the Malawi National AIDS Commission. The exact annual rate implied by the Wawer results is 12%. The Hollingsworth, Anderson, and Fraser (2008) reanalysis of the Wawer et al. (2005) data finds an annual transmission rate of 10.6% from asymptomatic partners (HIV-positive sex partners who have not just recently contracted the virus and do not yet have AIDS), which are the majority of cases, but does not provide an overall average.

[25] The questions respondents asked were recorded on the baseline survey.

to white. Enumerators then counted and circled these transmissions.

D.1 Information Treatment Oral Script and Visuals

[Read the text in this script to the respondent. Do not show it to them, and do not show them the pictures until instructed.]

Now I'm going to tell you about some recent research on HIV in Africa that you may not have heard about. People usually think that if they are married to someone who is HIV positive they are sure to be positive themselves. Have you heard a man say "I don't need to get tested, my wife got tested at the hospital so I know I'm the same as she is"? Sometimes people say the same thing about casual partners. Have you heard people say "If you lie together you die together"?

This study was about couples where one partner contracted HIV, and the researchers wanted to see if the other partner would also become HIV-positive.

- The researchers studied about 100 couples in Uganda. In each couple, one partner had HIV and one did not.
- All of the couples were having sex without using condoms. Most of the couples had sex about once every three days.
- In this picture, the black people represent someone in the couple who is HIV-negative while the white people represent their HIV-positive sex partners.

[Show the respondent the first picture.
Explain that there are 100 couples shown, and what the colors mean.]

- Remember, there were 100 people at the beginning who did not have HIV, and some of those 100 people contracted HIV.
- One year later, after all the couples were having sex without condoms, how many of the 100 uninfected people do you think got HIV?

Number:_____

[Show the first picture again, and circle a group of couples equal to the number the respondent chose.]

- Actually, one year later, the researchers came back and tested those people, and only about 10 of the partners had contracted HIV

[Show the respondent the second picture.
Count the 10 new white stick figures.
Circle all 10 couples with two white partners, and then show the first picture again to demonstrate the difference between what the respondent chose and the actual number.]

- In the picture, just 10 of the 100 black partners – the people who initially did not have HIV – has turned white.
- Remember, all of these people were having sex with someone who was HIV-positive. Most of them did not get HIV.
- This means that if someone has sex with an HIV-positive person without a condom, they may not necessarily contract HIV themselves.
- Even though this research was in Uganda, the Malawi National AIDS Commission, NAC, has found that the same thing is true here in Malawi.

What do you think about this? Could this be true? [ ] Yes [ ] No

Why/why not? _____

What do you believe about this research? _____

[Check responses against IT1 to IT9. If they match any of the options, read those answers and tick the boxes]

1

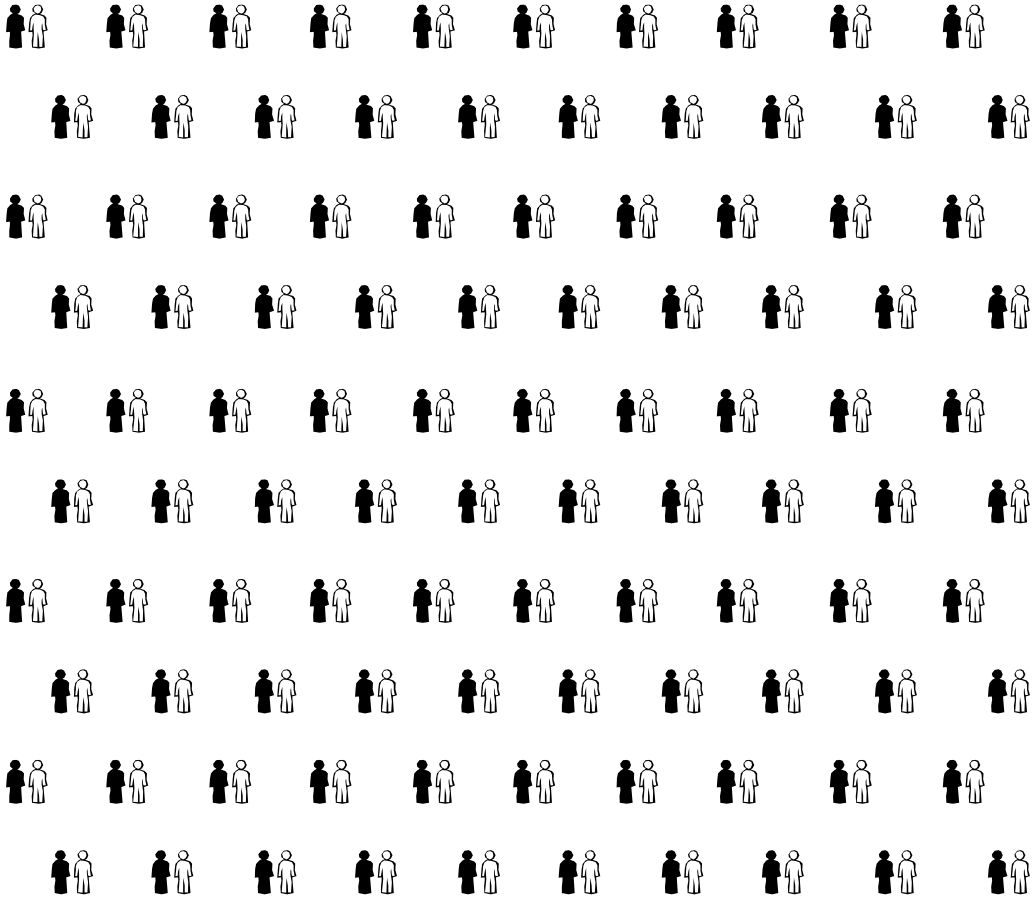

2

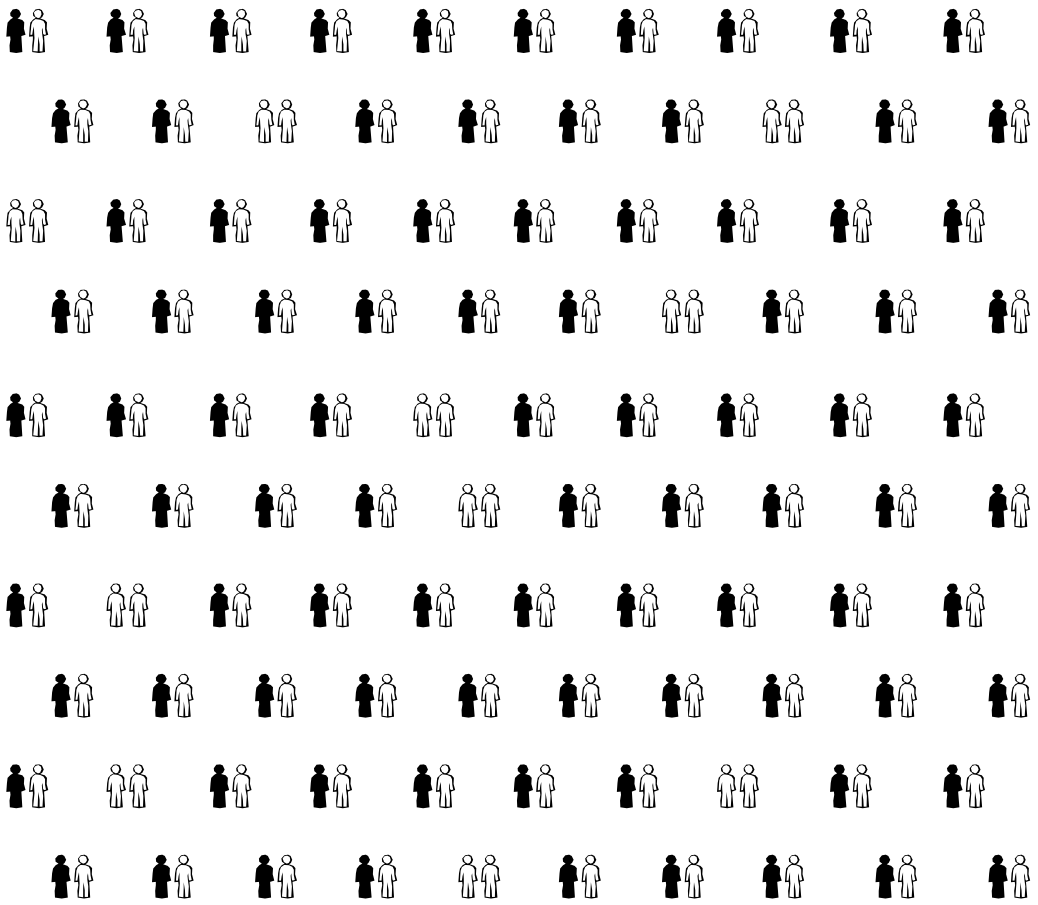

This is true because:

- People contract HIV from exposure to infected fluids (e.g. semen, vaginal fluid and blood). But sometimes the virus can't get into the body, even if exposed fluids touch someone's vagina or penis.
- And sometimes, when a person is exposed to infected fluids through sex, their body can fight it off and keep them from catching it.
- The amount of the virus in an infected person's body also varies. Sometimes it's more and sometimes it's less, and the more there is of the virus, the more likely a person's sex partner is to become infected.
- Also, having sex with too little vaginal fluid, and too much friction, can increase the risk, especially if there is bleeding. But you can still get HIV even without friction or bleeding.

Does this mean that people can't contract HIV from an infected partner?

- No! Some people who had sex with an HIV-positive partner did get HIV, but not all of them did.
- Having unprotected sex with an HIV-positive person is very dangerous, but it is not certain to infect you every time
    - In fact, almost everyone who contracts HIV in Malawi gets it through vaginal sex. But it doesn't happen for sure, just because you have sex with an infected partner one time.

The longer people had an infected sex partner for, the more likely they were to contract HIV.

- The more times you have unprotected sex with an HIV-positive person, the higher your chance of contracting HIV becomes.
- If you have sex with an HIV-positive person very few times, your chance of contracting the virus is small.
- If you have sex with an HIV-positive person many times, your chance of contracting the virus is large.
- Think about the Yabooka contests that AIRTel has, or the Tikolore contests from TNM. Some people win airtime, money, or a car, but others don't. Again, if you play just once, you aren't likely to win. But if you play a lot of times, your chances of winning improve.

[Show the respondent both pictures again, and emphasize the difference between the number circled on the first page and the true number, 10, on the second page.]

Do you have any questions?

[Standardized responses to questions and statements. <u>DO NOT READ THESE UNLESS THE RESPONDENT ASKS ABOUT THEM</u>. For each type of question that is asked, tick the appropriate box above the question. <u>Be sure to read the responses and tick the boxes for any reasons the respondent gave above under "why/why not?"</u>

[ ] IT1
<u>If respondent says "half-half"/"theka-theka" or "half can get it, half cannot get it"/"theka litha kutenga, theka litha osatenga":</u>
A: Yes, some can some can not, but more not than yes.

[ ] IT2
<u>USE THIS ANSWER FOR ANY FOLLOWUP QUESTION ABOUT "WHY"</u>
Q: Why do people sometimes get HIV from unprotected sex and sometimes not get it?
A: Every time you have unprotected sex with an infected partner, there's a chance you will get the virus. Why do people sometimes win the Airtel or TNM game and sometimes lose? It's just a chance.

[ ] IT3
Q: How is it possible for someone to have sex with an HIV-positive person many times and not get HIV?
A: Sometimes it's possible for people to get lucky, even if they have unprotected sex with an infected partner many times. How is it possible to play the Airtel game for many weeks and not win? It's just a chance.

[ ] IT4
Q: Is this because people with blood group O are immune to HIV?
A: No. People with all blood groups have equal chances to get HIV from unprotected sex.

[ ] IT5
Q: Is this because people can only get HIV from someone with the same blood group?
A: No. People with any blood group can get HIV from someone with any other blood group.

[ ] IT6
Q: Is this because some of the people had sex with less friction and more fluid?
A: No. Sex with less friction and more fluid is safer, but you can still get HIV from unprotected sex even if there is less friction and more fluid.

[ ] IT7
Q: Is this because some of the people had bleeding during sex and some didn't?
A: No. Bleeding makes sex more dangerous, but you can get HIV from unprotected sex even if there is no bleeding.

[ ] IT8
Q: Is this because there was a mistake or the testing equipment failed?
A: No. The researchers confirmed the tests by triple-checking all of them with different testing equipment.

[ ] IT9
Q: Is this because some (or all) of the people used condoms when they had sex?
A: No. None of the people used condoms when they had sex.

[ ] IT10
ALL OTHER QUESTIONS:
A: I can't provide any information on that topic.
IF RESPONDENT DOES ASK OTHER QUESTIONS, RESPOND AS ABOVE AND DESCRIBE BRIEFLY HERE:

_____

_____ **]**

E Proof that Controlling for Baseline Values of the Outcome Variable Minimizes the Bias in Estimated Treatment Effects

Consider estimating the effect of a randomly assigned treatment T on outcome y. The typical econometric strategy for analyzing experiments is to estimate

$$y_i^e = \alpha + \beta_{POST} T_i + e_i \tag{E1}$$

That is, regress endline values of the outcome on an indicator for treatment status plus a constant. $\hat{\beta}_{POST}$ will consistently estimate the causal effect of T on y due to the random assignment of the treatment. When baseline data is available, it is also common to use difference-in-difference specifications which utilize first differences of the outcome and treatment status as the dependent and independent variable respectively:

$$Dy_i = \alpha + \beta_{DIFF} DT_i + e_i \tag{E2}$$

Here $Dy_i \equiv y_i^e - y_i^b$ and $DT_i \equiv T_i^e - T_i^b = T_i$, and β_{DIFF} also consistently estimates the parameter of interest. Frison and Pocock (1992) show that both β_{POST} and β_{DIFF} have higher variance than a third alternative, which includes baseline values of the outcome of interest as a control in a regression of endline outcomes on treatment status:[26]

$$y_i^e = \alpha + \beta T_i + \gamma y_i^b + e_i \tag{E3}$$

$\hat{\beta}$ is also consistent for the effect of T on y; as it is more efficient, it is preferable on those grounds alone. However, $\hat{\beta}$ has a further advantage in the case of (even slight) baseline imbalance in an outcome variable: it is also less biased than either other option.

Let $d^b = \bar{y}_T^b - \bar{y}_C^b$ be the baseline difference in the outcome of interest, and σ^2 be the

[26] This is also referred to as the "ANCOVA" (analysis of covariance) estimator in the medical literature, where the relevant alternatives were variants of analysis of variance ("ANOVA") methods.

variance of the error term. The variance of the error can be decomposed into a component due to measurement error (σ_e^2), and a remaining component $\sigma^2 - \sigma_e^2$. Frison and Pocock (1992) show that for a single baseline and followup the bias due to baseline imbalance is given by:

1. $Bias_{POST} = \frac{\sigma^2\rho}{\sigma^2-\sigma_e^2}d^b$ for the POST estimator,
2. $Bias_{DIFF} = \frac{\sigma^2(\rho-1)+\sigma_e^2}{\sigma^2-\sigma_e^2}d^b$ for the DIFF estimator, or
3. $Bias_{OPTIMAL} = \frac{\sigma_e^2\rho}{\sigma^2-\sigma_e^2}d^b$ for the optimal estimator.

It is important to note that although the size of the bias term will diminish as d^b falls, it will be nonzero unless d^b is identically zero. Thus these finite-sample bias terms are potentially relevant even if the outcome is balanced in the sense of not having statistically significant differences at baseline. Frison and Pocock show that the relative size of $Bias_{POST}$ and $Bias_{DIFF}$ depends on whether ρ is greater or less than 0.5, and note that in most cases σ_e^2 will be very small relative to $\sigma^2 - \sigma_e^2$ so that $Bias_{OPTIMAL}$ is nearly zero. However, it is also possible to show the intuitive result that, in addition to having lower variance than the alternatives, $\hat{\beta}$ is also uniformly less biased in the presence of baseline imbalance in a finite sample. Consider the relative size of the bias terms,

$$\frac{Bias_{DIFF}}{Bias_{OPTIMAL}} = \frac{\sigma^2(\rho-1)+\sigma_e^2}{\sigma^2-\sigma_e^2}\frac{\sigma^2-\sigma_e^2}{\sigma_e^2\rho} = \frac{\sigma^2(\rho-1)+\sigma_e^2}{\sigma_e^2\rho} \tag{E4}$$

And

$$\frac{Bias_{POST}}{Bias_{OPTIMAL}} = \frac{\sigma^2\rho}{\sigma^2-\sigma_e^2}\frac{\sigma^2-\sigma_e^2}{\sigma_e^2\rho} = \frac{\sigma^2}{\sigma_e^2} \tag{E5}$$

Each of these ratios approaches infinity as the portion of variance due to measurement error approaches zero, and reaches a minimum value of 1 if $\sigma_e^2 = \sigma^2$. This is equivalent to saying that 100% of the residual variance of y is due to measurement error; we can rule that out in the case of sexual activity since our regression model will logically predict only a

small portion of the true variation in patterns of sex. Thus, when the baseline mean of the outcome of interest is not identical across the treatment and control groups, $\hat{\beta}$ will be less biased than $\hat{\beta}_{POST}$ or $\hat{\beta}_{DIFF}$.

This derivation is confirmed by a simple simulation of the DGP described above. Appendix Figure E1 shows the results of simulating the DGP 1000 times and computing the bias of each estimator. The green squares show the binned average of estimates from the optimal estimator, while the red diamonds show the binned average bias for the DIFF estimator and the blue circles show the binned average bias for the POST estimator. The optimal estimator's bias always lies between that of the DIFF and POST estimators, and in expectation it is less than that of the other two estimators when the treatment-control difference is not zero.

Appendix Figure E1
Bias of Different Estimators as a Function of the
Baseline Treatment-Control Difference in Outcomes

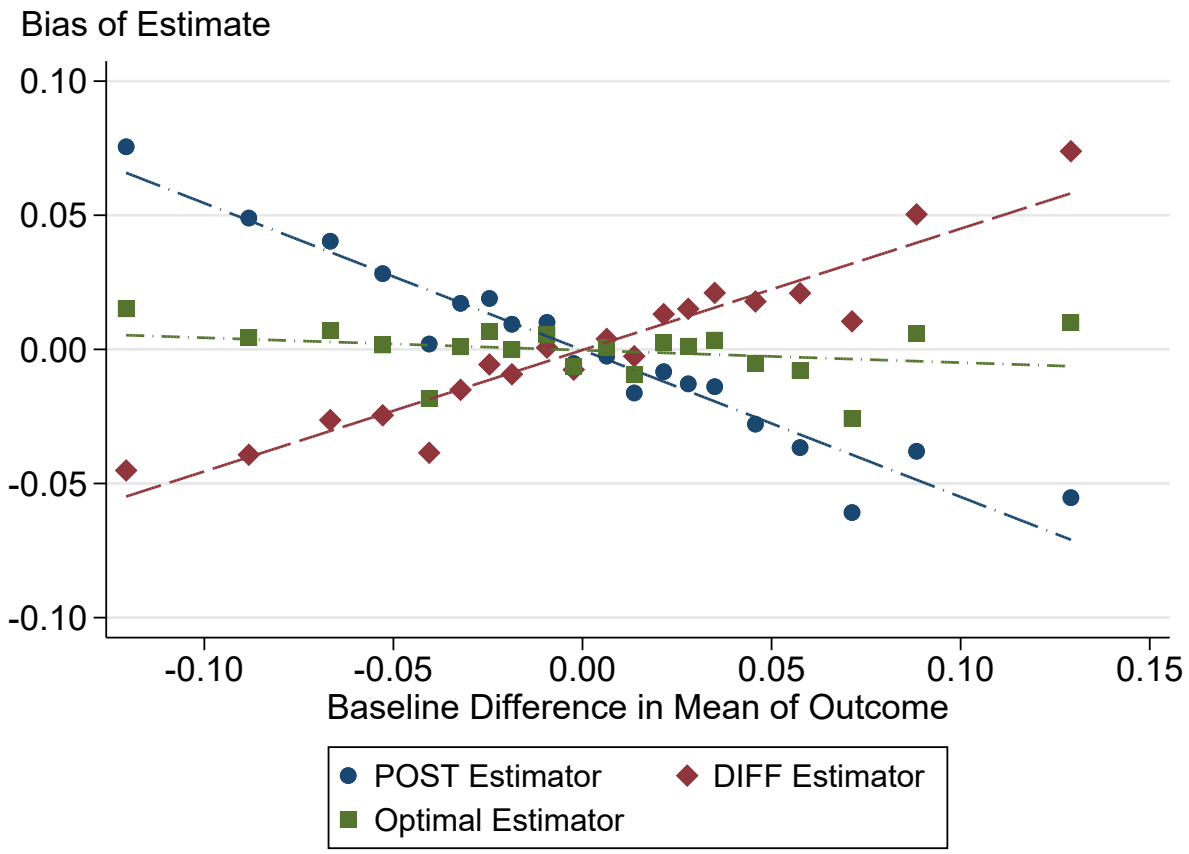

F Average Treatment Effects and 2SLS Estimates

I estimate the average effect of the information treatment on sexual behavior using the following regression.

$$y_i = \beta_0 + \beta_1 T_i + \lambda y_i^b + Z_i'\eta + \varepsilon_i \tag{F1}$$

The impact of the treatment on sexual activity is small in magnitude: it is possible to rule out changes larger in magnitude than 20 percent. The number of sex acts in the past week

rises by 10 percentage points. Focusing specifically on the margin of abstinence (whether people have any sex at all), this shifts by 5 percentage points, which is roughly 0.1 standard deviations. The risk indices confirm that these results are robust to multiple hypothesis testing: both the overall and sex diary risk indices rise by 6%, significant at the 10% and the 5% level respectively. The treatment has no effect on condom use, nor on condom purchases. This is consistent with the extremely high rates of unprotected sex: at baseline just 1 in 10 sex acts involved a condom, leaving limited room for increases in risk-taking at this margin.

Appendix Table F1
Average Treatment Effects

	Any Sex in Past Week	Log Sex Acts in Past Week	Log Unprotected Sex Acts in Past Week	Log Sex Acts in Past 30 Days	Log Sex Partners in Past 30 Days	Log Condoms Acquired in Past 30 Days	Log Overall Sexual Activity Index
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treatment Group	0.050**	0.101**	0.071	0.057	0.012	0.080	0.077*
	(0.024)	(0.047)	(0.045)	(0.058)	(0.019)	(0.075)	(0.041)
Observations	1,292	1,292	1,292	1,271	1,290	1,283	1,261
Adjusted R-squared	0.238	0.277	0.260	0.346	0.288	0.140	0.388
Control-group Mean	0.490	1.673	1.481	5.339	0.767	2.523	-0.0258
Control-group SD	0.500	2.385	2.286	6.382	0.576	9.658	0.994

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. All regressions include controls for sampling strata and baseline values of the outcome variable; Column 6 was not measured at baseline so the baseline values for Column 5 are used as a proxy. Heteroskedasticity-robust standard errors, clustered by village, in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

The randomized information treatment was an exogenous shock that could only have affected endline sexual activity through its effect on risk beliefs. This allows me to estimate the risk elasticity of sexual activity via 2SLS, as follows:

$$x_i^e = \beta_0 + \beta_1 T_i + \lambda y_i^b + \mu x_i^b + Z_i'\eta + \varepsilon_i \tag{F2}$$

$$y_i^e = \beta_0 + \beta_1 x_i^e + \lambda y_i^b + \mu x_i^b + Z_i'\eta + \varepsilon_i \tag{F3}$$

where x_i^e and x_i^b are endline and baseline risk beliefs respectively, and likewise for y_i^e and y_i^b; all other variables are defined as in Section 2.

The 2SLS and OLS estimates are shown in Panels A and B of Appendix Table F2 respectively. The OLS regressions are estimated on the control group only. The OLS results have a uniform positive bias relative to 2SLS, confirming that OLS is not consistent in this context. This concords with the results in Oster (2012), who finds that OLS estimates of the elasticity of sexual behavior with respect to the true prevalence of HIV are biased and wrong-signed.

The bias of the OLS estimates implies that the omitted variable in the second-stage regression is positively correlated with risk beliefs. There are at least two potential explanations for this pattern. The first is reverse causality due to endogenous information acquisition: sexual activity may directly drive risk beliefs rather than vice versa. For example, people who have more risky sex may decide as a result of their high number exposures to seek out information about HIV risks. As HIV risk messaging typically overstates how easy the virus is to contract, this could lead to higher risk beliefs.[27] Second, some other variable could drive both sexual activity and HIV risk beliefs. One such possible factor is sociability: people who are more sociable are likely to have more sex partners and also be exposed to more gossip about HIV, which would tend to replicate the common messaging that HIV is extremely easy to get. Either of these patterns implies that people with high risk beliefs—whom my main results show are at risk of fatalism—are also, in the *status quo*, those who have more risky sex. This could mean that fatalism is even more important for

[27] A similar empirical pattern is documented in Gerrard et al. (1996).

public health policy than the size of the fatalistic group would imply, because some research suggests that HIV epidemics are predominantly driven by a small group of people who have high levels of sexual activity (Koopman, Simon, and Riolo 2005).

Appendix Table F2

Comparison of 2SLS and OLS Estimates of the Effect of Endline Risk Beliefs on Sexual Activity

	Any Sex in Past Week	Log Sex Acts in Past Week	Log Unprotected Sex Acts in Past Week	Log Sex Acts in Past 30 Days	Log Sex Partners in Past 30 Days	Log Condoms Acquired in Past 30 Days	Log Overall Sexual Activity Index
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: 2SLS Estimates							
Endline Risk Belief	-0.249**	-0.532**	-0.360	-0.236	-0.037	-0.380	-0.347*
	(0.123)	(0.242)	(0.229)	(0.280)	(0.098)	(0.393)	(0.198)
Observations	1,268	1,268	1,268	1,248	1,266	1,259	1,238
Adjusted R-squared	0.206	0.252	0.251	0.335	0.279	0.129	0.367
1st-Stage F-Statistic	184.2	184.1	184.9	190.5	186	185.8	191.2
Control-group Mean	0.492	1.674	1.487	5.396	0.775	2.574	-0.0164
Control-group SD	0.500	2.387	2.290	6.415	0.575	9.751	0.992
Panel B: OLS Estimates (Control Group Only)							
Endline Risk Belief	0.189***	0.331***	0.264**	0.755***	0.222***	0.185	0.512***
	(0.058)	(0.116)	(0.109)	(0.204)	(0.065)	(0.193)	(0.136)
Observations	632	632	632	624	632	632	624
Adjusted R-squared	0.208	0.196	0.155	0.215	0.159	0.010	0.232
Control-group Mean	0.492	1.674	1.487	5.396	0.774	2.570	-0.0203
Control-group SD	0.500	2.387	2.290	6.415	0.575	9.744	0.993

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. All regressions include controls for sampling strata and baseline values of the outcome variable; Column 6 was not measured at baseline so the baseline values for Column 5 are used as a proxy. 2SLS estimates use the randomized treatment group assignment as an instrumental variable for endline risk beliefs. OLS estimates use the endline data for the control group only. Heteroskedasticity-robust standard errors, clustered by village, in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

G Robustness Checks

My results are robust to a wide range of robustness checks. This Appendix begins with an overview of the findings; this is then followed by the detailed results.

G.1 Overview of robustness tests.

G.1.1 Variations in the Measure of Sexual Activity

I begin by showing that the pattern of fatalism—negative treatment effects for the highest risk beliefs—is not specific to my main outcome variable. Figure G1 presents five alternative measures of sexual activity: any sex in the past week (Panel A), unprotected sex acts in the past week (Panel B), sex acts in the past 30 days (Panel C), sex partners in the past 30 days (Panel D), and the combined sexual activity index (Panel E).

The same basic pattern of heterogeneity is evident in these outcome variables as well, albeit more-noisily in Panels C and D. The effects in Panel A show that my results are not driven by outliers, since in that specification the outcome is binary. The statistically significant fatalism effect for the overall index in Panel E ($p = 0.049$) suggests that my results do not depend on the specific outcome I use. A related robustness check is presented in Appendix G.2, which shows that my findings are not sensitive to using the unlogged versions of the outcome variables. This provides further reassurance that outliers are not driving my results.

G.1.2 Variations in the Definition of Risk Beliefs

The belief measure I use in this study is the same one that I used for the empirical analysis in Kerwin (2012). In that paper, I document a U-shaped relationship between sexual activity and risk beliefs in data from another part of southern Malawi. I chose the specific definition of risk beliefs—the per-act transmission rate from unprotected sex times the local prevalence of HIV among attractive people—to capture the risk of having sex with a random potential sex partner. In Appendix subsubsection G.6.1, I show that his choice does matter: when I break out the two components of risk beliefs, the pattern of fatalism is driven by the prevalence beliefs rather than the transmission rate beliefs (Appendix Figure G6). However,

the latter distribution has a large mass point at 100%: nearly half of the sample thinks that a single exposure to HIV will certainly lead to an infection. My measure of risk beliefs breaks up that mass point by how likely an unprotected sex act is to lead to an HIV exposure, that is, the prevalence of HIV.

Moreover, my results are robust to a number of other potential definitions of risk beliefs. First, they are qualitatively unchanged if I use annual, rather than per-act, risk beliefs (Appendix Figure G9). In Appendix G.6.3, I show that my findings are also robust to correcting for the interviewer knowledge spillovers documented in Kerwin and Ordaz Reynoso (2021) and to controlling for baseline interviewer fixed effects.

Rather than the risk of contracting HIV from a *random* sex partner, people may react to the risk from their *current* partner. In Appendix G.6.4, I show that people in committed relationships exhibit the same pattern of fatalism if I replace the local prevalence of HIV with the perceived likelihood that their primary sex partner has HIV.

My results are also not materially affected by the way I handle initial responses of 50 percent (Appendix G.6.5). As described in Section 1.2, I interpret responses of 50 percent as potentially indicating uncertainty rather than the respondent's actual beliefs; thus enumerators were instructed to follow up and ask for the respondent's best guess. I use those best guesses (rather than the initial response of 50 percent) in my analysis. This decision does not matter for my results: the treatment effects do not differ substantially for people who changed their responses. Moreover, the pattern of fatalism is visible even if I control for the interaction between changing one's response and the treatment indicator.

G.1.3 Changes in the Regression Specification

My results are also robust to interacting a linear term in baseline risk beliefs with the treatment, rather than the decile-based approach in my main specification. This approach uses the following specification:

$$y_i = \beta_0 + \beta_1 T_i + \beta_2 x_i + \beta_3 T_i \times x_i + \lambda y_i^b + \mu T_i \times y_i^b + \sum_{j=1}^{J} \left[\gamma_j w_i^j + \delta_j T_i \times w_i^j\right] + Z_i'\eta + \varepsilon_i \tag{G1}$$

Panel A of Table G1 presents estimates of Equation G1. Column 1 shows that, on average, the treatment increases sexual activity in the past seven days by 10 log points, or 11 percent. Thus the average treatment effect is consistent with standard risk compensation. The estimates of Equation G1 in column 2 show that the treatment effect varies substantially by baseline risk beliefs. For a baseline risk belief of 100 percent, the treatment effect is a statistically significant *decline* of 36 log points, or 30 percent. This decline is statistically significant at the 0.05 level. This pattern is consistent with fatalism: the treatment convinces people with high risk beliefs that they are not doomed to HIV infection, raising the marginal cost of risky sex. Adding controls for interactions between the treatment and baseline covariates does not substantively affect this result (column 3).

In addition, the same qualitative pattern emerges if I divide baseline risk beliefs using 5, 15, or 20 quantiles of baseline beliefs, instead of using deciles (Figure G2). The treatment effect for the top bracket of baseline risk beliefs is significant at the 10 percent level if I use 5 quantiles and at the 5 percent level if I use 15 or 20.

G.1.4 Corrections for Potential Confounders in Baseline Risk Beliefs

Since the baseline risk belief variable that is at the core of my identification strategy is not randomly assigned, it is possible that other factors correlated with risk beliefs could drive my results. I address this in several ways. My main specification controls for interactions between the treatment indicator and every baseline balance variable from Table A1. In Appendix Figure G5, I show that my main results are robust to omitting all of the controls from the regression. Relatedly, the fatalism results pass "placebo tests" where I put the baseline value of the outcome variable on the left-hand side of the regression (Appendix G.3).

Another potential threat to identification is that the treatment and control groups are

slightly imbalanced on sexual activity measures *within* the top decile of risk beliefs (Table 1, Panel B). To address this imbalance on observables, I re-estimate my main treatment effects using only people from this sample. The results, in Appendix Table G3, confirm the same qualitative pattern from my main specification in Panel B of Figure 1. Column 2 includes controls for all of the covariates in the table and their interactions with the treatment indicator. This makes the estimated treatment effect even larger, and it remains statistically significant. The specific magnitude should be interpreted with caution, however: this regression includes just 54 villages, and uses 44 of those degrees of freedom.

A closely connected concern is that *unobserved* factors that are correlated with risk beliefs could be leading to the negative treatment effects at the top of the risk belief distribution. This parallels the concern that Altonji, Elder, and Taber (2005) address for average treatment effects. These unobserved confounders would have to be correlated with both high risk beliefs and responses to the information treatment. One example of a potential confounder is the propensity to believe information one is told about HIV, which could lead to higher risk beliefs and also to more updating of one's beliefs in response to the treatment. To account for this possibility, I show that my results are robust to Oster (2019) bounds. Specifically, show that selection on unobservables would have to be 2.5 times as strong as selection on observables to explain the estimated fatalism effects (Appendix G.5.)

G.1.5 Additional Robustness Checks

I also conduct a number of other robustness checks. The same pattern of treatment effect heterogeneity holds for both men and women (Appendix Figure G14). Thus my results cannot be explained by gender differences in risk beliefs, and are not specific to an arbitrary subset of the population.

As part of the information intervention, enumerators asked respondents if they believed the information and provided scripted responses to common questions and concerns. Beliefs update substantially even for respondents who said they did not initially believe the information, indicating that the enumerators were successful in addressing their doubts. Consistent with this, the pattern of fatalism is visible both for people who did and did not initially believe the information script (Appendix G.8).

Basic knowledge of how HIV is transmitted is high. The survey contained a set of questions about whether various activities can spread HIV, such as blood transfusions (yes) and sharing food (no). In Appendix G.9, I show that there are no large differences across study arm or between fatalistic and non-fatalistic in terms of answers to these questions. I also show that controlling for the interactions of these questions with the treatment indicator does not change my main results on fatalism.

Appendix Figure G1

Robustness to Alternate Outcome Variables

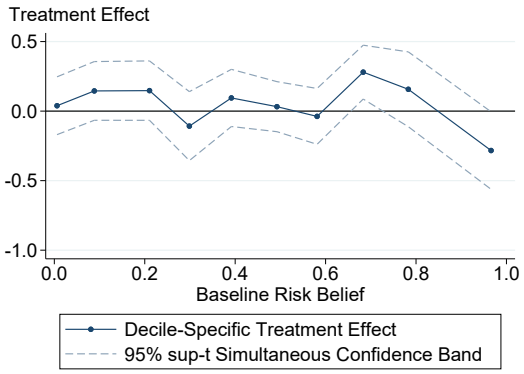

Panel A: Any Sex in Past Week

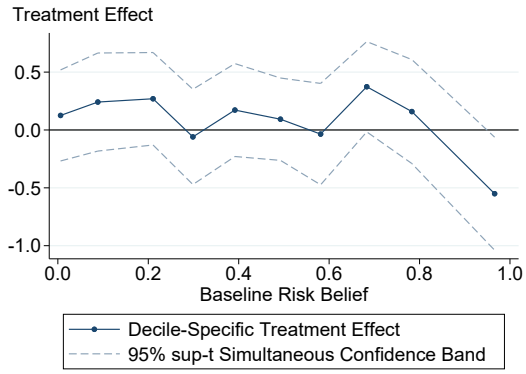

Panel B: Ln(Unprotected Sex Acts in Past Week)

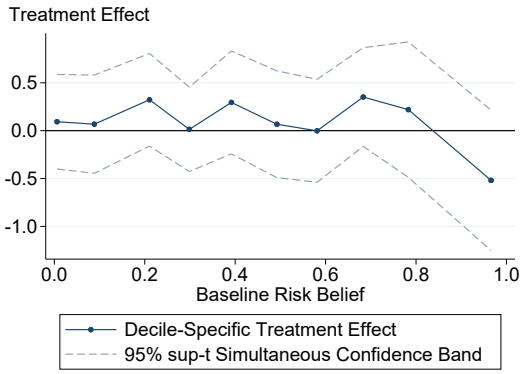

Panel C: Ln(Sex Acts in Past 30 Days)

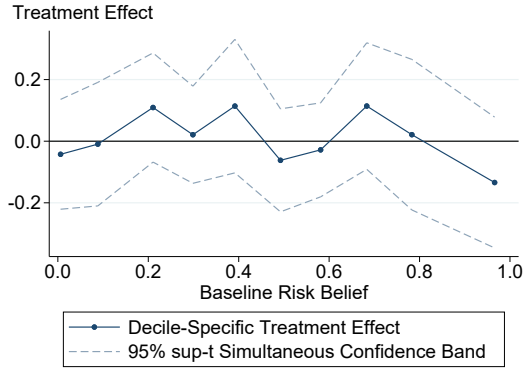

Panel D: Ln(Sex Partners in Past 30 Days)

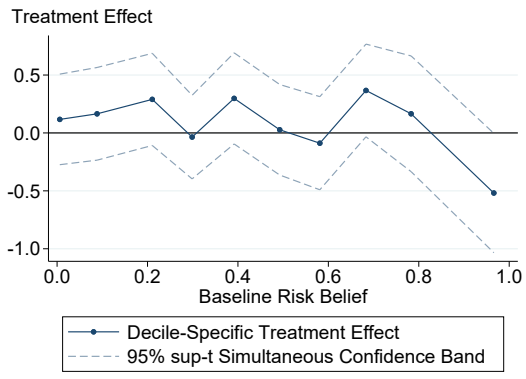

Panel E: Ln(Overall Sexual Activity Index)

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.

Appendix Figure G2
Treatment Effect Heterogeneity by Quantiles of Baseline Risk Beliefs
Robustness to Varying Number of Brackets of Beliefs
Outcome: Ln(Sex Acts in Past 7 Days)

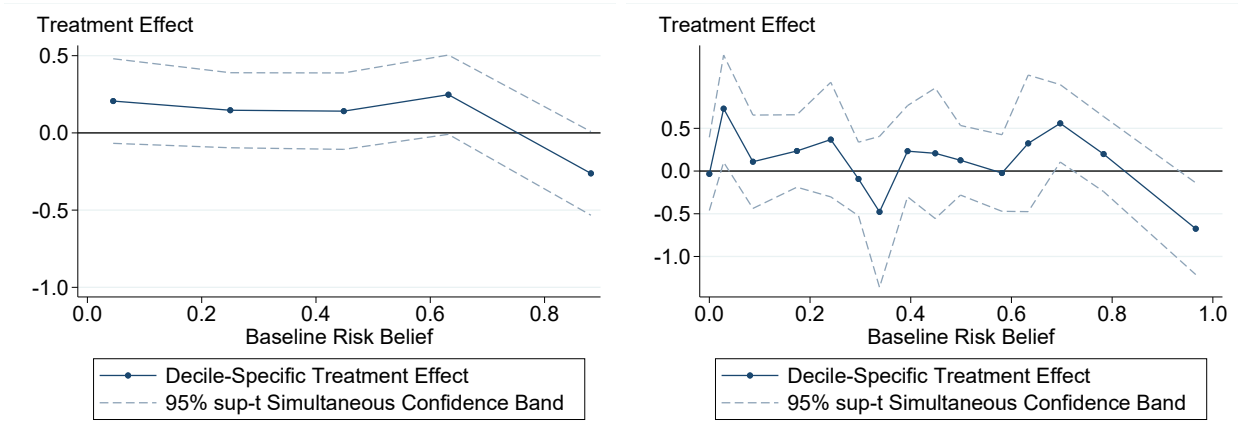

Panel A: 5 Brackets of Beliefs instead of 10

Panel B: 15 Brackets of Beliefs instead of 10

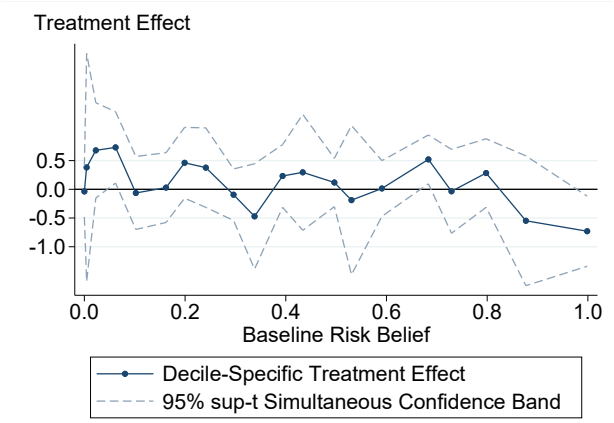

Panel C: 20 Brackets of Beliefs instead of 10

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each quantile.

Appendix Table G1

Treatment Effects on Sexual Activity by Baseline Risk Beliefs, Linear Model

	Outcome: Log Sex Acts in Past Week								
	Panel A: Main Specification			**Panel B:** No Control for Baseline Outcome			**Panel C:** No Controls		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treatment (T)	0.101**	0.116**	0.110**	0.055	0.076	0.102**	0.055	0.079	0.102**
	(0.047)	(0.046)	(0.049)	(0.058)	(0.057)	(0.048)	(0.061)	(0.060)	(0.051)
T*(Baseline Risk Belief [0-1])		-0.477***	-0.426**		-0.449**	-0.411**		-0.441**	-0.431***
		(0.162)	(0.169)		(0.192)	(0.166)		(0.190)	(0.161)
Control for Baseline (BL) Outcome	Yes	Yes	Yes	No	No	No	No	No	No
Stratification Cell FEs	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
T Interacted w/BL Outcome	No	No	Yes	No	No	No	No	No	No
T Interacted w/Other BL Covariates	No	No	Yes	No	No	Yes	No	No	Yes
Observations	1,292	1,275	1,232	1,292	1,275	1,232	1,292	1,275	1,232
Adjusted R-squared	0.277	0.284	0.297	0.016	0.023	0.287	0.000	0.008	0.286
Control-group Mean	1.673	1.673	1.683	1.673	1.673	1.683	1.673	1.673	1.683
Control-group SD	2.385	2.382	2.390	2.385	2.382	2.390	2.385	2.382	2.390
Treatment Effect for BL Belief=1		-0.360**	-0.315*		-0.373*	-0.309*		-0.362*	-0.329*
		(0.173)	(0.178)		(0.202)	(0.176)		(0.204)	(0.174)

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). All regressions include controls for sampling strata and baseline values of the outcome variable. Main effects are included for all variables included in interactions. Other baseline covariates include the complete set of variables included in Table A1. Heteroskedasticity-robust standard errors, clustered by village, in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

G.2 Unlogged outcomes

Appendix Figure G3 shows that my main results also hold if I do not log the outcome variable. Panel A shows my main outcome variable, sex acts in the past week, while Panels B through E show the other continuous outcome variables from Figure G1. (Panel A of Figure G1 shows any sex in the past week, which is discrete and thus was not logged). The same pattern is evident in the unlogged specifications as in the logged versions, with a large negative treatment effect on sexual activity for the top decile. This effect is statistically significant at the 0.1 level for the overall outcome index in Panel E ($p = 0.033$).

Appendix Figure G3
Robustness to Unlogged Outcome Variables

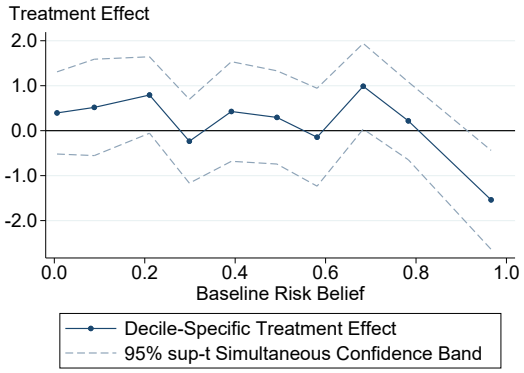

Panel A: Sex Acts in Past Week

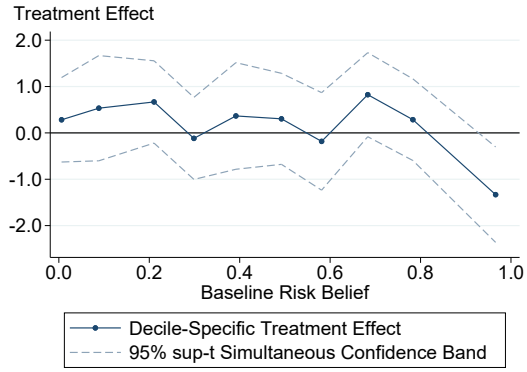

Panel B: Unprotected Sex Acts in Past Week

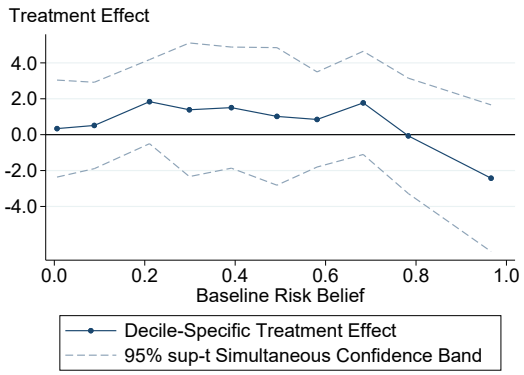

Panel C: Sex Acts in Past 30 Days

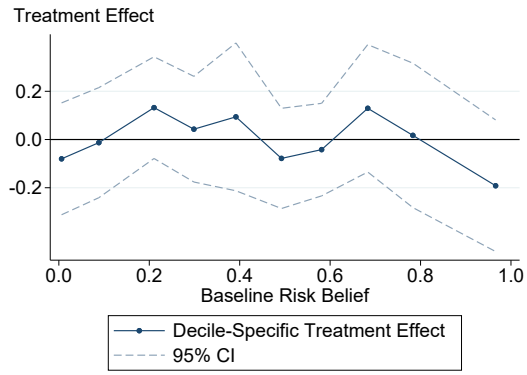

Panel D: Sex Partners in Past 30 Days

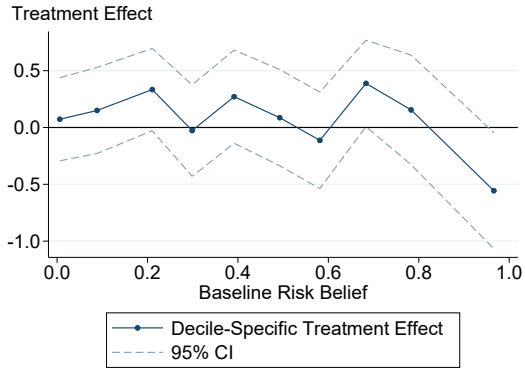

Panel E: Overall Sexual Activity Index

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.

G.3 Balance

Another potential concern is balance: are the treatment and control group balanced on baseline covariates not just overall, but also within the fatalistic subset of respondents? To further explore balance I run "placebo" regressions, where the outcome is levels of sexual activity measured at baseline rather than endline. To avoid controlling for outcomes in these regressions, I alter the controls by omitting the baseline value of the outcome variable as well as all the main effects and interactions for the sexual activity variables from Table A1.

Appendix Table G2 shows the linear specification (Equation G1). The interaction between the treatment indicator and baseline risk beliefs is statistically insignificant regardless of whether I control for sampling strata fixed effects, and the sign of the interaction coefficient is positive, rather than negative as in the results for the actual outcome variable in Table G1. Appendix Figure G4 presents the non-linear specification from Equation 1. There are no statistically significant treatment effects at any decile.

Appendix Figure G4
Treatment Effect Heterogeneity by Deciles of Baseline Risk Beliefs
Placebo Test
Outcome: Ln(Sex Acts in Past 7 Days) at Baseline

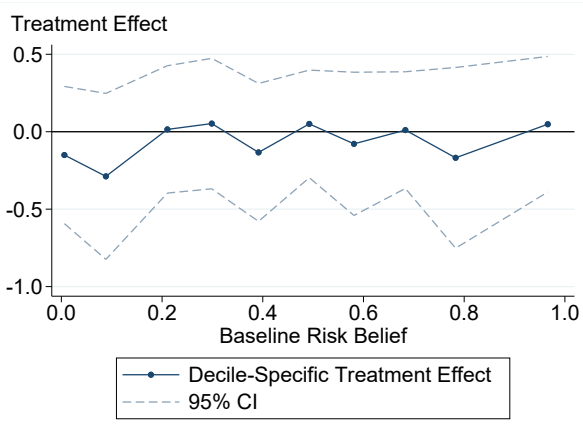

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.

Appendix Table G2
Treatment Effect Heterogeneity by Baseline Risk Beliefs
Placebo Test (Outcome Measured at Baseline)

	Outcome: Log Sex Acts in Past Week (Baseline)					
	Panel A: Main Specification			Panel B: No Controls		
	(1)	(2)	(3)	(7)	(8)	(9)
Treatment (T)	-0.089*	-0.077	-0.050	-0.086*	-0.074	-0.048
	(0.049)	(0.049)	(0.045)	(0.050)	(0.051)	(0.046)
T*(Baseline Risk Belief [0-1])		0.052	0.175		0.064	0.185
		(0.177)	(0.168)		(0.182)	(0.172)
Control for BL Outcome	No	No	No	No	No	No
Stratification Cell FEs	Yes	Yes	Yes	No	No	No
T Interacted w/BL Outcome	No	No	Yes	No	No	No
T Interacted w/Other BL Covariates	No	No	Yes	No	No	Yes
Observations	1,292	1,275	1,255	1,292	1,275	1,255
Adjusted R-squared	0.005	0.009	0.143	0.001	0.006	0.143
Control-group Mean	0.237	0.244	0.245	0.237	0.244	0.245
Control-group SD	0.976	0.977	0.978	0.976	0.977	0.978
Treatment Effect for BL Belief=1		-0.025	0.125		-0.010	0.138
		(0.185)	(0.176)		(0.190)	(0.180)

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Main effects are included for all variables included in interactions. Other baseline covariates include the complete set of demographic variables included in the second section of Table A1. Heteroskedasticity-robust standard errors, clustered by village, in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Consistent with these placebo tests, Appendix Figure G5 shows that my main results are robust to keeping just the stratification cell fixed effects (Panel A) and to dropping all the controls from the regression (Panel B). Panels B and C of Table G1 show that the linear heterogeneity results are also robust to dropping the controls for the baseline outcome variable and the stratification cells.

Appendix Figure G5

Treatment Effect Heterogeneity by Deciles of Baseline Risk Beliefs
Robustness to Omitting Controls
Outcome: Ln(Sex Acts in Past 7 Days)

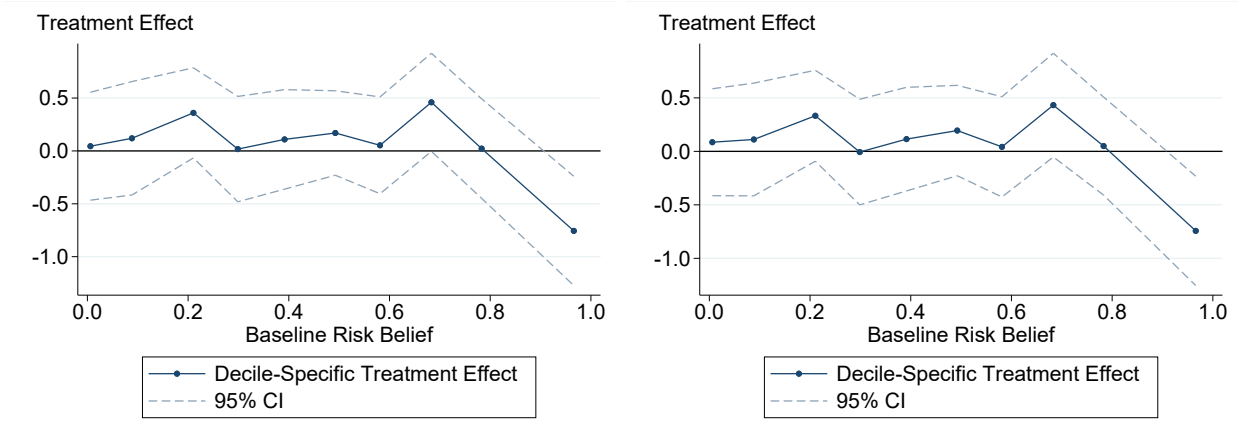

Panel A: Controlling for Sample Strata Only

Panel B: No Controls

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.

G.4 Potential Imbalance within the Top Decile of Risk Beliefs

Appendix Table G3 shows treatment effect estimates specifically for the top decile of baseline risk beliefs.

Appendix Table G3
Treatment Effects on Sexual Activity,
Restricting Sample to Top Decile of Baseline Risk Beliefs

	Outcome: Log Sex Acts in Past Week	
	(1)	(2)
Treatment (T)	-0.690***	-1.122***
	(0.174)	(0.222)
Control for Baseline (BL) Outcome	Yes	Yes
Stratification Cell FEs	Yes	Yes
T Interacted w/BL Outcome	No	Yes
T Interacted w/Other BL Covariates	No	Yes
Observations	112	106
Adjusted R-squared	0.305	0.242
Control-group Mean	2.594	2.632
Control-group SD	2.475	2.473

Notes: Sample includes 112 people from 57 villages for whom both baseline and endline surveys were successfully completed, and who were in the top decile of baseline risk beliefs. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) × (Baseline Prevalence Belief). All regressions include controls for sampling strata and baseline values of the outcome variable. Main effects are included for all variables included in interactions. Other baseline covariates include the complete set of variables included in Table A1, with the exception of the Muslim indicator which is dropped due to collinearity. Heteroskedasticity-robust standard errors, clustered by village, in parentheses.Randomization inference p-values, adjusted for sampling strata and clustered by village: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

G.5 Oster (2019) bounds

Another potential issue is that the fatalistic people might differ on some *unobserved* variable that drives treatment effect heterogeneity. Since people's risk beliefs are not randomly assigned, I cannot depend on randomization to guarantee balance (in expectation) on unobserved covariates. Oster (2019) proposes a test for the degree of unobserved selection that

would be needed in order to drive an estimated treatment effect, along the same lines as Altonji, Elder, and Taber (2005). The test is based on assuming that selection on unobserved variables follows a similar pattern to selection on observed variables, which in this case are the interactions between the treatment indicator and baseline covariates, $T_i \times w_i^j$. Omitting them from my regression reduces the magnitude of the treatment effect for people in the top decile of risk beliefs by just 2%—from 0.662 to 0.648—and shifts the R-Squared from 0.338 to 0.346. To compute a bound on how strong selection on unobservables would have to be than selection on observables, I assume that including all unobservables would raise the R-squared to 1.0. Under that assumption, I find selection on unobservables would need to be nearly two and a half times as strong as that on observables to explain away my results.

G.6 Definitions of baseline Risk Beliefs

G.6.1 Composite Risk Beliefs

I break up the main risk belief variable into its two components—per-act transmission rate beliefs and prevalence beliefs—in Appendix Figure G6. In Panel A, the treatment effects are significantly lower for the top decile of transmission rate beliefs than the first decile ($p = 0.065$); this is inconsistent with a basic model of risk compensation, which would predict larger effects for the higher risk belief category. However, there is no evidence of fatalism when examining the per-act risk alone. In contrast, Panel B shows statistically significant fatalistic risk responses among people with the highest prevalence beliefs.

One explanation for the lack of clear-cut fatalism in Panel A comes from the fact that the two belief variables are positively correlated with each other, but not strongly so—the Pearson correlation coefficient is 0.14. Appendix Figure G7 shows a binned scatterplot of prevalence beliefs against transmission risk beliefs, along with a line of best fit.

In particular, the correlation is one-sided: there is more variation in transmission beliefs for people with high prevalence beliefs than vice versa. Many of people who think the transmission rate is 100% believe the prevalence of the virus is quite low, so their effective risk from unprotected sex is not particularly high. Panel A of Appendix Figure G8 shows that the median person in that group thinks the local prevalence of the virus is 50%. People with high prevalence beliefs, on the other hand, almost all think the transmission rate is high as well. The average per-act risk belief for the top decile of prevalence beliefs is 87%, and nearly two thirds of the people in that group think the transmission rate is 100% (Panel B). Thus, on average, the perceived chance of contracting HIV from a single sex act is just 50% for people who think the transmission rate is 100%, but is nearly 100% for people who think the prevalence is 100%.

Appendix Figure G6

Treatment Effect Heterogeneity by Deciles of Baseline Risk Beliefs
Showing Components Separately
Outcome: Ln(Sex Acts in Past 7 Days)

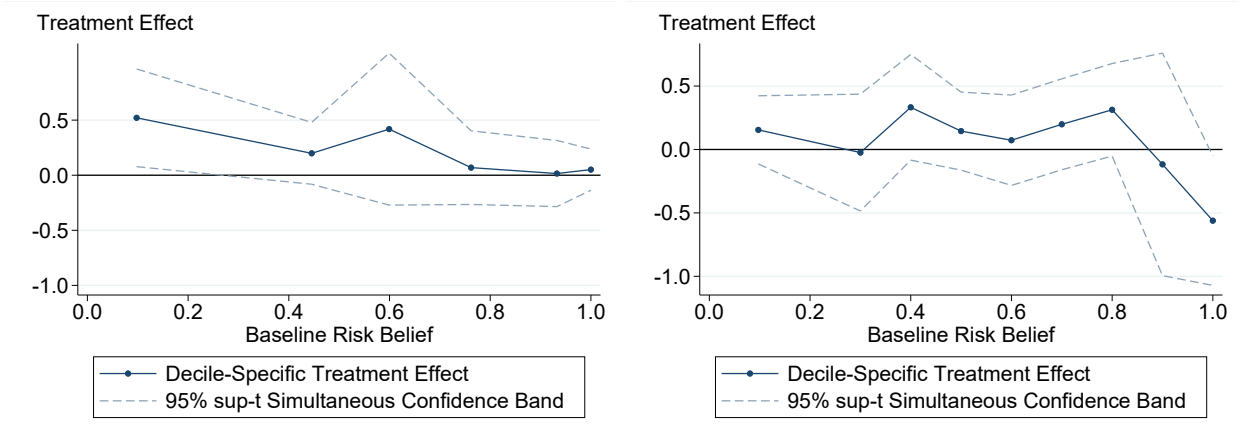

Panel A: Beliefs: Per-Act Transmission Rate

Panel B: Beliefs: Prevalence of HIV

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each quantile.

Appendix Figure G7

Binned Scatterplot of Prevalence Beliefs vs. Transmission Risk Beliefs

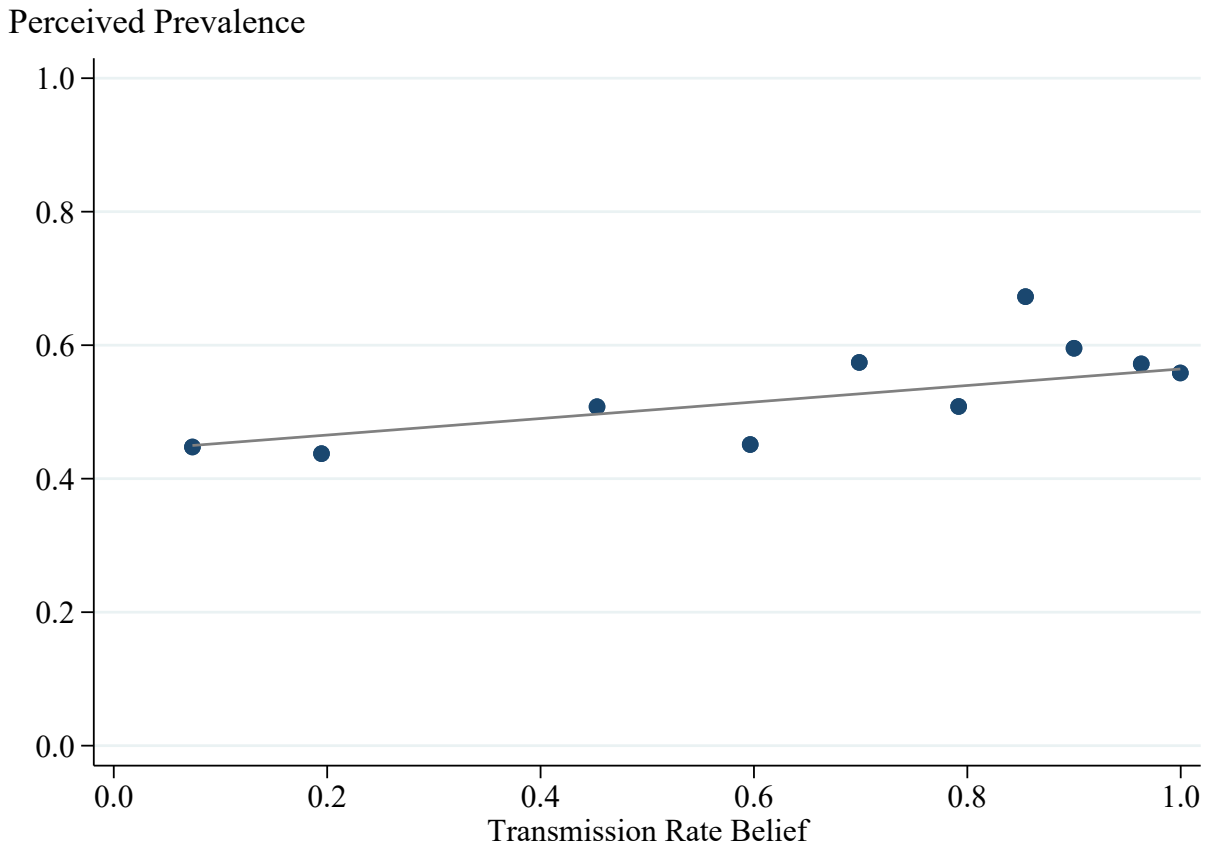

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed.

Appendix Figure G8
Histograms of Baseline Risk Belief Components
For People with High Beliefs on Other Component

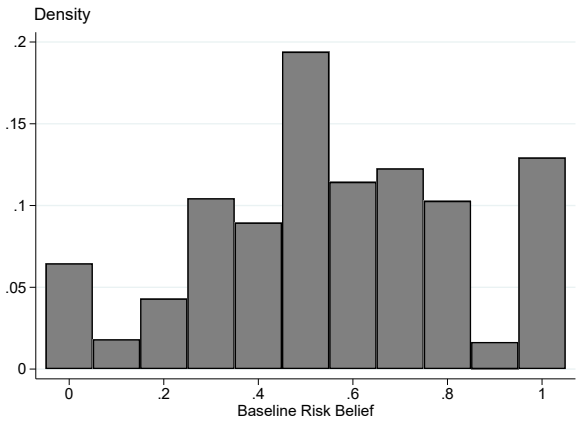

Panel A: Prevalence Beliefs for People Who Believe Transmission Rate is 100%

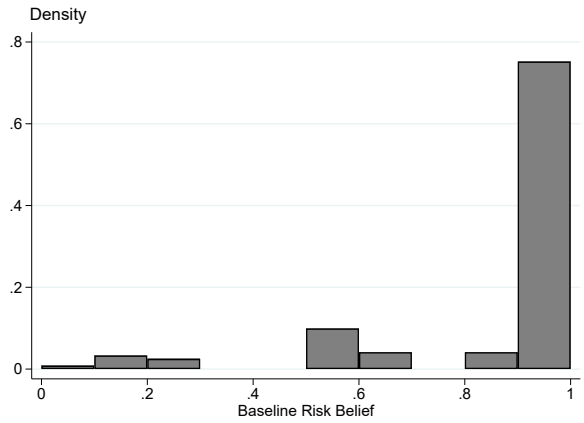

Panel B: Transmission Rate Beliefs for People Who Believe Prevalence is 100%

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed.

G.6.2 Annual Risk Beliefs

Motivated by my previous work on the topic (Kerwin 2012), my main risk belief variable is the per-act risk of HIV transmission from a random attractive person from the local area. The information treatment, however, taught people about annual risks instead, because they are easier to explain. It is thus important to assess whether my results are robust to using the annual risk instead of the per-act one. The answer is yes: Appendix Figure G9 shows that the same basic pattern of treatment effect heterogeneity is visible if I use annual risks rather than per-act risks, for both endline risk beliefs (Panel A) and sexual activity (Panel B).

Appendix Figure G9
Treatment Effect Heterogeneity by Deciles of Baseline Risk Beliefs
Annual Risk Beliefs

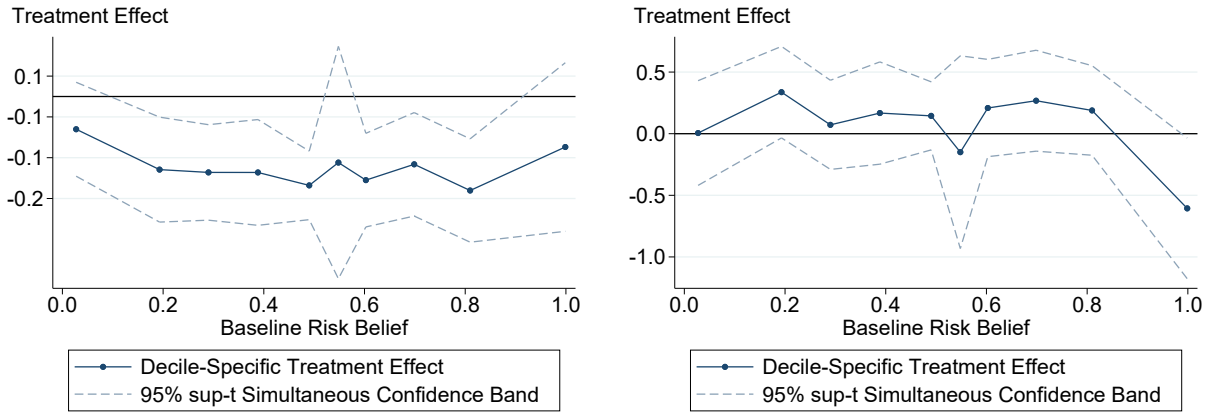

Panel A: Endline Risk Beliefs

Panel B: Ln(Sex Acts in Past 7 Days)

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from one year of regular unprotected sex with a randomly chosen attractive person of the opposite sex from the local area. The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.

G.6.3 Interviewer knowledge spillovers onto measured baseline risk beliefs

As noted in Section 1, the measured values of baseline risk beliefs are slightly imbalanced due to interviewer knowledge spillovers (Kerwin and Ordaz Reynoso 2021). The interviewer knowledge effects are shown in Appendix Figure G10; the baseline difference in risk beliefs

is much smaller than the actual treatment effect on endline risk beliefs. Interviewer effects are also visible in the treatment-group distribution of endline risk beliefs, which are strongly dependent on who the baseline interviewer was. The interviewer-specific means vary from 0.04 to 0.25, and a cluster-adjusted F-test easily rejects joint equality ($F(11, 34) = 39.25$, $p < 0.001$).

Appendix Figure G11 addresses this issue in two ways. Panel A adjusts the beliefs via linear regression, subtracting off separate linear time trends within each study arm as well as the estimated trend break. Specifically, I estimate:

$$x_i^b = \beta_0 + \beta_1 Date + \beta_2 Post + \beta_3 Post \times Date + \varepsilon_i \tag{G2}$$

and then construct

$$\tilde{x}_i^b = x_i^b - \hat{\beta}_1 Date - \hat{\beta}_2 Post - \hat{\beta}_3 Post \times Date \tag{G3}$$

which is equivalent to subtracting off the slopes of the first two lines shown in Appendix Figure G10 as well as the level difference between them; I then run my main specification using $\tilde{x}_i^b$ instead of x_i^b. Panel B of Appendix Figure G11 computes the deciles of baseline risk beliefs within each study arm, rather than across both, eliminating the effect of any shifts in risk beliefs. Panel C uses the original belief variable, but controls for baseline interviewer fixed effects. The same pattern of fatalism from Panel B of Figure 1 is visible in all three panels.

Appendix Figure G10

Measured Risk Beliefs over Time, by Study Arm

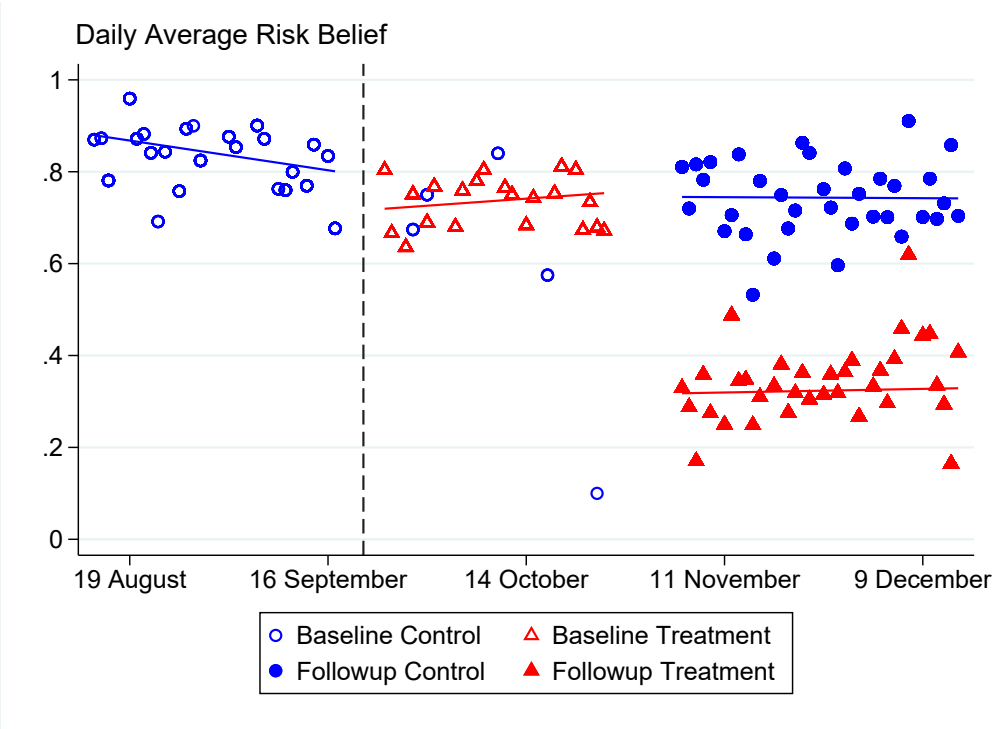

Notes: Sample includes 1,292 sexually active adults who were successfully interviewed at both baseline and endline. Risk beliefs are the perceived probability of contracting HIV from a single unprotected sex act with an infected partner. Each point represents the mean value of the risk beliefs for a given day; baseline control beliefs are hollow circles, endline control beliefs are solid circles, baseline treatment beliefs are hollow triangles, and endline treatment beliefs are solid triangles. The lines are linear fits of beliefs on date for a given date range and study arm. The dashed vertical line indicates the date of the training sessions when the survey interviewers were trained to provide the information treatment about HIV transmission risks.

Appendix Figure G11
Treatment Effect Heterogeneity by Deciles of Baseline Risk Beliefs
Robustness to Adjusting Beliefs
Outcome: Ln(Sex Acts in Past 7 Days)

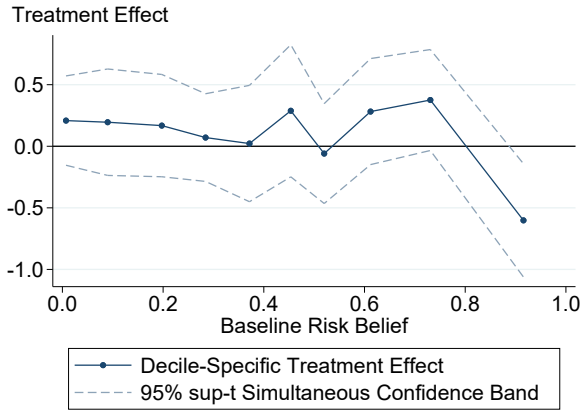

Panel A: Adjusting Beliefs for Linear Trends w/a Break

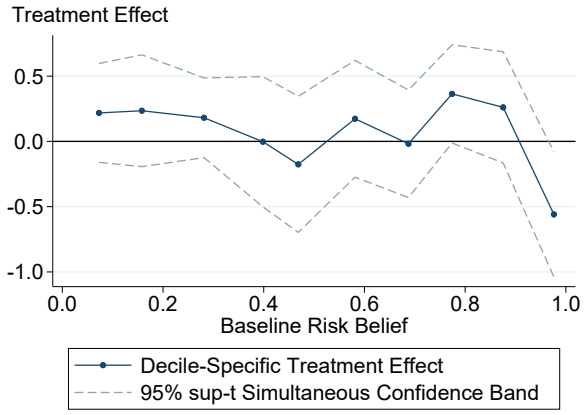

Panel B: Using Within-Arm Percentiles of Beliefs

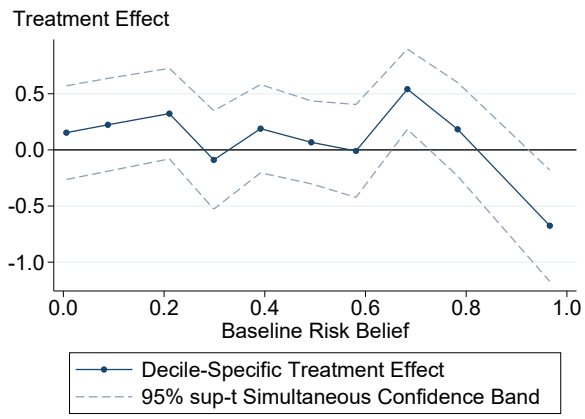

Panel C: Controlling for Enumerator Fixed Effects

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.

G.6.4 Beliefs about Risk from Primary Sex Partner

My main risk belief variable captures "community risk"—the risk of sexual activity with a random person from the local community. An alternate definition of risk beliefs is to interact the perceived transmission rate with beliefs about the HIV status of one's primary sex partner, which I will refer to as "partner risk". The survey was designed to not explicitly ask whether this was the respondent's spouse, in order to encourage honest responses in a context where infidelity is common.[28] Given high rates of infidelity, risk beliefs based on the baseline primary sex partner are measured with some amount of error: some respondents have exclusively had sex with one person in the past month and do not suspect infidelity, but would consider having sex with other people in the future. Reflecting this fact, at least 15% of people appear to change partners between waves of the survey based on differences in the reported length of the relationship. The rate of changing partners is only slightly lower for people who were initially married. As a result we would expect the beliefs about community risks to apply to some of the people who are coded as being in committed relationships, whereas the beliefs about partner risks should apply only to those who are in committed relationships.

The data reflects exactly that pattern. Appendix Figure G12 presents treatment effect heterogeneity by community risks (my main risk belief variable) and partner risks (swapping prevalence for the partner's HIV status).[29] These are broken out by whether the respondent is in a committed relationship, defined as one where they do not suspect their partner of infidelity and have not had any other sex partners in the past 30 days. Panels A and C show that people outside committed relationships react fatalistically only to community risks and not to partner risks, while Panels B and D show that people who are in committed relationships react fatalistically to both types of risk (sup-t adjusted p-values = and 0.05 and 0.07 respectively). I can reject equal treatment effects for the top quantile of community

[28] One fifth of respondents suspect their partner of cheating, and 3% of respondents admit to cheating themselves. The rate of self-reported infidelity was even higher (4.5%) for people who volunteered (unprompted) that their primary sex partner was their spouse. Unfaithfulness in marital relationships in southern Malawi has been documented in extensive previous research (Schatz 2005, Conroy 2014), including infidelity by married women (Tawfik and Watkins 2007).

[29] These results are for total sex acts in the past week. Because the survey was designed to not capture the exact identity of sex partners, I am unable to determine the number of sex acts with the primary sex partner from baseline.

and partner risks for people who are not in committed relationships ($p = 0.004$), but not for those in committed relationships ($p = 0.290$).[30]

Appendix Figure G12
Treatment Effect Heterogeneity by Deciles of Baseline Risk Beliefs
Community Risk vs. Partner Risk
Outcome: Ln(Sex Acts in Past 7 Days)

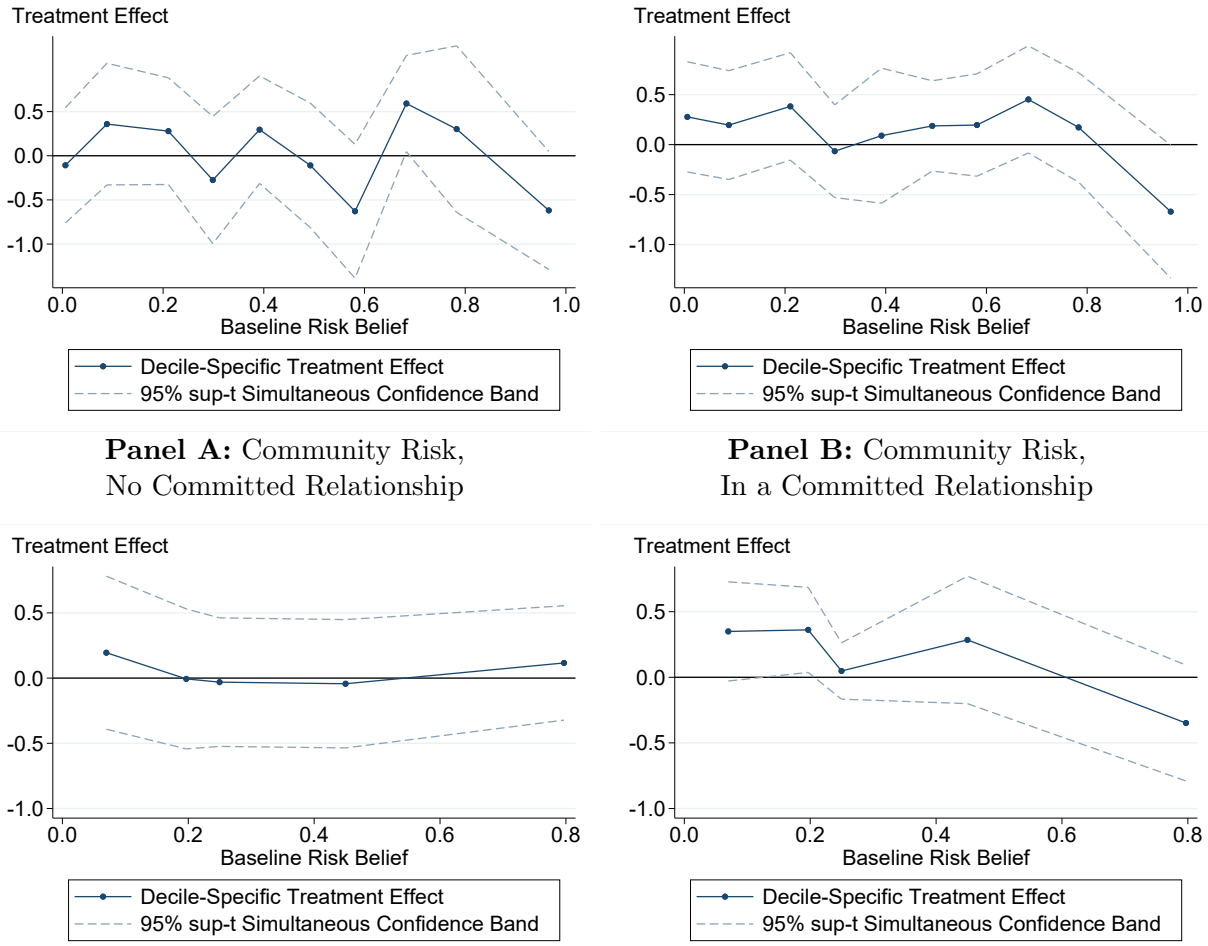

Panel A: Community Risk, No Committed Relationship

Panel B: Community Risk, In a Committed Relationship

Panel C: Partner Risk, No Committed Relationship

Panel D: Partner Risk, In a Committed Relationship

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. "Community risk" is the product of the per-act transmission rate belief and the perceived local prevalence among attractive people. "Partner risk" is the product of the per-act transmission rate belief and the likelihood that one's primary sex partner had HIV at baseline. The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each quantile.

[30] These latter two p-values come from seemingly unrelated regression analyses and thus are not sup-t adjusted.

G.6.5 Initial Risk Belief Responses of 50%

As described in Section 1, responses of 50% to probability questions sometimes mean the respondent was simply unsure about the risk in question, rather than an actual belief that the risk is 50-50. I handle this by building on work by Hudomiet, Kézdi, and Willis (2011), who ask respondents in that category if they really think the answer is 50% or if they are just not sure. People who say they are just unsure are then asked for their best guess. At baseline, 34% of people give an initial answer of 50% to at least one of the two components of the risk belief variable, but just 4% are ambiguous about both answers. Out of those who give an answer of 50% on at least one of the two components, 26% revise their answer when given an opportunity to (or 9% of the entire sample).

People who initially give responses of 50%, or who update their answer when given a chance to, may respond differently to the information treatment. Since being in one of these two groups is correlated with baseline risk beliefs,[31] I explore differences in treatment effects for them using modified versions of my non-linear specification. Building on Equation 1, I add main effects and interactions with the treatment for indicators for being in each of the two groups (initial answer of 50% or changed response when given option). This examines whether there is an additional difference for people in either of those two groups, after allowing treatment effects to differ by the level of baseline risk beliefs. The results are shown in Appendix Table G4. There is no evidence of heterogeneity in either the updating of beliefs or the effects of the treatment on endline sexual behavior. Moreover, my main results are unchanged by the addition of these variables to the regression model.

[31] The average risk beliefs of people who initially answered 50% are 11 percentage points lower than those of the rest of the sample ($p < 0.001$), which in the expected direction given the very high average responses on the risk belief questions. The average risk beliefs of people who revise their answers are 3 percentage points lower than those of the rest of the population, but this difference is statistically insignificant ($p = 0.334$).

Appendix Table G4

Differences in Treatment Effects for People with Initial Risk Beliefs of 50% and Those Who Changed their Responses

	Endline Risk Beliefs (1)	Log Sex Acts in Past Week (2)	Endline Risk Beliefs (3)	Log Sex Acts in Past Week (4)
Treatment [T] X				
1st Decile of Baseline Beliefs	-0.050	0.153	-0.052	0.163
	(0.038)	(0.143)	(0.037)	(0.144)
2nd Decile of Baseline Beliefs	-0.129***	0.215	-0.128***	0.238
	(0.043)	(0.176)	(0.038)	(0.154)
3rd Decile of Baseline Beliefs	-0.092**	0.304**	-0.102**	0.319**
	(0.042)	(0.140)	(0.042)	(0.142)
4th Decile of Baseline Beliefs	-0.277***	-0.122	-0.277***	-0.119
	(0.042)	(0.157)	(0.043)	(0.152)
5th Decile of Baseline Beliefs	-0.203***	0.164	-0.203***	0.182
	(0.043)	(0.154)	(0.041)	(0.147)
6th Decile of Baseline Beliefs	-0.142***	0.041	-0.145***	0.092
	(0.051)	(0.182)	(0.038)	(0.131)
7th Decile of Baseline Beliefs	-0.303***	-0.031	-0.304***	-0.034
	(0.044)	(0.154)	(0.044)	(0.151)
8th Decile of Baseline Beliefs	-0.205***	0.512***	-0.205***	0.511***
	(0.043)	(0.141)	(0.043)	(0.138)
9th Decile of Baseline Beliefs	-0.244***	0.187	-0.246***	0.185
	(0.059)	(0.153)	(0.061)	(0.153)
10th Decile of Baseline Beliefs	-0.162**	-0.668***	-0.162**	-0.668***
	(0.077)	(0.183)	(0.077)	(0.184)
T X (Initially Answered 50%)	0.000	0.065		
	(0.030)	(0.128)		
T X (Changed Response from 50%)			-0.002	0.078
			(0.049)	(0.194)
Control for BL Outcome	Yes	Yes	Yes	Yes
Stratification Cell FEs	Yes	Yes	Yes	Yes
T Interacted w/BL Outcome	Yes	Yes	Yes	Yes
T Interacted with Other Baseline Covariates	Yes	Yes	Yes	Yes
Observations	1,212	1,232	1,212	1,232
Adjusted R-squared	0.216	0.306	0.215	0.306
Control-group Mean	0.352	0.176	0.352	0.176
Control-group SD	0.268	0.980	0.268	0.980

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. All regressions include controls for sampling strata and baseline values of the outcome variable. Heteroskedasticity-robust standard errors, clustered by village, in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

G.7 Gender differences in beliefs

A related issue has to do with differences in beliefs by gender. Appendix Figure G13 shows that men tend to have higher risk beliefs than women, and women are more likely to be in the top decile of risk beliefs. Thus estimated heterogeneity in treatment effects by risk beliefs could simply be picking up heterogeneity by gender. My main specification addresses this by including controls for both an indicator for being male and its interaction with the treatment indicator. To further explore this possibility, Appendix Figure G14 estimates Equation 1 separately by gender, showing that both men and women exhibit fatalism. The effects are stronger for men than for women. *A priori*, it is not clear whether we would expect this pattern or the opposite. On the one hand, men often have more agency in relationships and thus more scope to adapt their behavior in response to the information treatment. On the other hand, because women have less agency, they may have more inevitable exposures to HIV and thus be more prone to fatalism.

Appendix Figure G13
Baseline Risk Beliefs by Gender

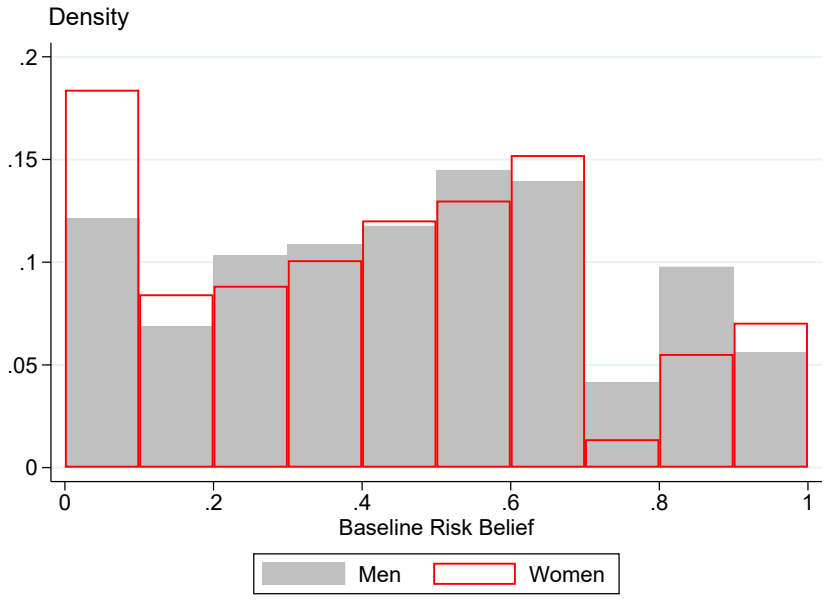

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) × (Baseline Prevalence Belief).

Appendix Figure G14
Treatment Effect Heterogeneity by Deciles of Baseline Risk Beliefs
and Gender
Outcome: Ln(Sex Acts in Past 7 Days)

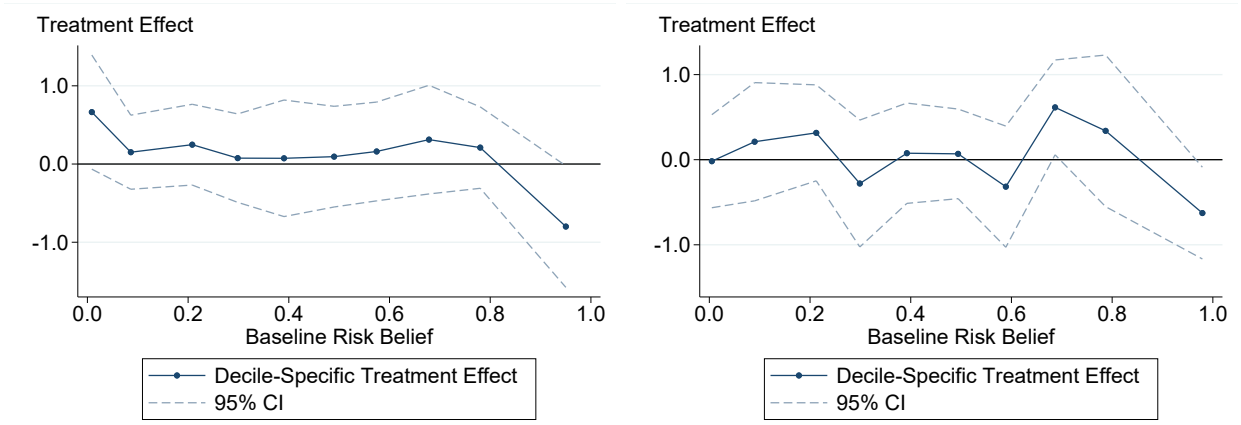

Panel A: Men

Panel B: Women

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.

In many contexts, rates of extramarital sex vary widely by gender, since most transactional sex workers are female. This could lead to floor effects in the female distribution, because it may not be possible to further reduce sexual activity from an already-low point. Indeed, on average women in my sample report 9 percent less sex than men do, although this difference is not statistically significant. However, there is little evidence of floor effects in the pattern of fatalism in Panel B of Appendix Figure G14. This can be explained by the fact that within the top decile of baseline risk beliefs, women actually report nearly 9 percent *more* sex than men do. One reason for the lack of a large gender gap in self-reported sexual activity in my sample is that in southern Malawi, transactional sex exists on a continuum, with women transitioning from sex workers to girlfriends to wives (Swidler and Watkins 2007). As a result, women who engage in transactional sex are more likely to show up in my sample than they would be in other settings.

G.8 Belief in the Information Treatment

After they were initially shown the information about HIV transmission risks, treatment-group respondents were asked if they believed what they were told, and why or why not. Interviewers were trained to answer several common reasons why people might not believe the risk information, but this still raises the question of whether people's initial inclination to believe the information altered how they responded to it. On this question, 39.1% of people initially said they did not believe the risk information. This helps explain why the treatment group's risk beliefs remain so high after the information treatment—the average person in the treatment group still thinks the annual risk of HIV transmission from unprotected sex with an infected partner is over 33%. People who initially disbelieve the information update their beliefs by 6.4 percentage points less than people who did believe it; average effects on sexual behavior are 2.7 percentage points smaller, but this difference is not statistically significant (Appendix Table G5). The same pattern of fatalism is visible for both groups (Appendix Figure G15), and there is no statistically significant difference in the decline in sexual activity for people in the top decile of risk beliefs.

Appendix Table G5

Treatment Effect Heterogeneity by Initial Belief in Information Treatment

	Outcome: Endline Risk Belief	Outcome: Log Sex Acts in Past Week
	(1)	(2)
Treatment Group	-0.207***	0.118**
	(0.015)	(0.058)
(Treatment Group) × (Don't Believe Information)	0.064***	-0.031
	(0.018)	(0.070)
Observations	1,249	1,289
Adjusted R-squared	0.206	0.278
Control-group Mean	0.351	0.170
Control-group SD	0.268	0.980

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. All regressions include controls for sampling strata and baseline values of the outcome variable. Heteroskedasticity-robust standard errors, clustered by village, in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Appendix Figure G15

Treatment Effect Heterogeneity by Deciles of of Baseline Risk Beliefs
and Initial Belief in Information Treatment
Outcome: Ln(Sex Acts in Past 7 Days)

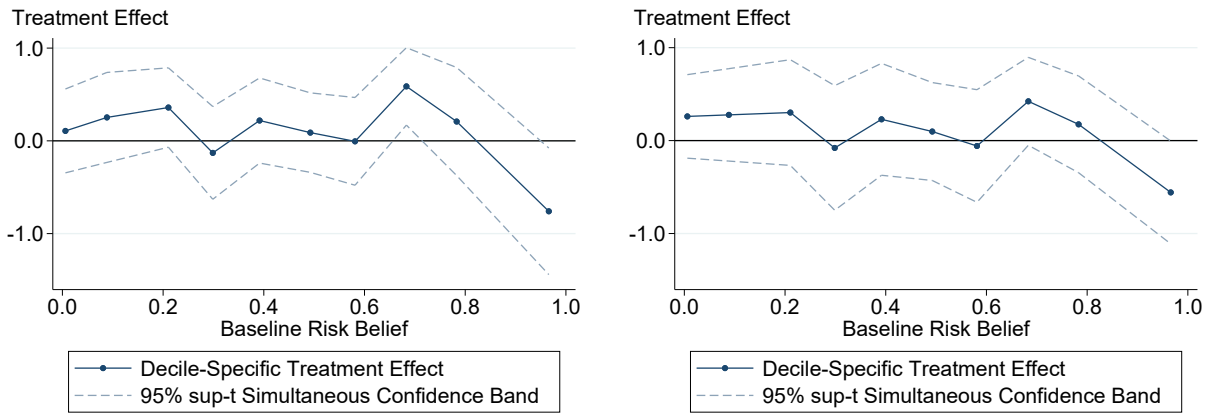

Panel A: Initially **Did** Believe Information **Panel B:** Initially Did **Not** Believe Information

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.

G.9 Basic Knowledge of HIV Risks

Both the treatment and control group were provided with written information about the basics of HIV prevention at the end of the baseline survey: that it is spread through vaginal sex and that condoms are effective at preventing it. (Note that this was separate from the information treatment, which was provided solely to the treatment group.) The specific text of the information that was provided is as follows:

> "Thank you very much for your time in taking this survey. I would now like to tell you some information about HIV prevention. According to the Malawi National AIDS Commission, Malawi is still experiencing a severe HIV epidemic. Most of the spread of the virus is through sex. Almost every person who has HIV now got it from having unprotected sex with an infected partner. However, not all sex is risky – if you have sex with a condom, even if your partner is infected, your chance of contracting HIV is very low. The National AIDS Commission therefore recommends that condoms be used whenever you are having sex with a partner who is HIV-positive, or whose HIV status you do not know. Condoms are a safe and effective means of preventing HIV transmission."

If there were baseline differences in HIV knowledge across study arms or between people with high risk beliefs and the rest of the sample, this could have led to differences in sexual activity across groups at endline. However, knowledge about HIV is extremely high in my sample. Based on the risk belief questions, 99.9% of respondents believed that HIV could be spread through unprotected sex, and 98.2% believe that condoms reduce that risk.

In addition to the questions about transmission probabilities, my data also contains a battery of questions about the ways in which people think HIV is transmitted. As another check on the levels of basic knowledge about HIV in my sample, and whether it is balanced by treatment status and baseline risk beliefs, Appendix Table G6 replicates Table A1 for these questions. Out of the seven options, the two correct answers are vaginal sex and blood transfusions. This data confirms that knowledge about HIV transmission is very high in my sample: roughly 90 percent of respondents think HIV can be transmitted via blood transfusions, and nearly 100 percent think it can be spread via vaginal sex. None of the

wrong answers are given even a third of the time. The sample is well-balanced on the HIV knowledge questions. There are no large differences in answers to the questions between the treatment and control groups; the treatment group is less likely to report that HIV can be spread by vaginal sex ($p = 0.038$) but the difference is just a single percentage point relative to a control-group mean of 99.7%.

When we compare the fatalistic and non-fatalistic samples, we see that the latter is 4 percentage points more likely to (correctly) report that it is not possible to contract HIV by sharing food ($p = 0.018$) relative to a control-group mean of 6.8%. None of the other differences in responses are large or statistically significant. To examine whether part of the observed pattern of fatalistic responses to the information treatment is driven by differences in HIV knowledge, I add these questions to the set of variables w_i in Equation 1, for which I include main effects and interactions with the treatment as controls. The results are shown in Figure G16. Including these additional controls has no appreciable effect on my results, which are nearly identical to those in Panel B of Figure 1.

Appendix Table G6
Balance for HIV Knowledge Questions

	Panel A: Control (C) vs. Treatment (T)				Panel B: Non-Fatalistic (N) vs. Fatalistic (F)			
	C	T	Diff. (p-value)	Obs.	N	F	Diff. (p-value)	Obs.
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Thinks HIV can be transmitted by								
Mosquitos	0.304	0.263	-0.042	1,291	0.285	0.279	-0.012	1,274
	(0.460)	(0.441)	(0.122)		(0.452)	(0.451)	(0.765)	
Shaking hands	0.068	0.065	-0.003	1,292	0.064	0.071	0.008	1,275
	(0.252)	(0.247)	(0.830)		(0.246)	(0.259)	(0.769)	
Vaginal sex	0.997	0.986	-0.011**	1,292	0.992	0.991	-0.002	1,275
	(0.056)	(0.117)	(0.038)		(0.088)	(0.094)	(0.821)	
Kissing	0.432	0.445	0.015	1,289	0.440	0.402	-0.037	1,272
	(0.496)	(0.497)	(0.643)		(0.497)	(0.492)	(0.460)	
Sharing food	0.054	0.074	0.021	1,290	0.068	0.027	-0.042**	1,273
	(0.227)	(0.263)	(0.176)		(0.252)	(0.162)	(0.018)	
Using the same toilet	0.076	0.095	0.019	1,287	0.087	0.071	-0.017	1,271
	(0.265)	(0.293)	(0.310)		(0.282)	(0.259)	(0.408)	
Blood transfusions	0.896	0.887	-0.008	1,290	0.891	0.893	0.002	1,273
	(0.306)	(0.317)	(0.645)		(0.312)	(0.311)	(0.934)	

Notes: Sample includes 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. p-values in Column 3 of each panel are adjusted for sampling strata and clustered by village: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Appendix Figure G16

Treatment Effect Heterogeneity by Deciles of Baseline Risk Beliefs
Robustness to Adding Controls & Interactions for HIV Knowledge
Outcome: Ln(Sex Acts in Past 7 Days)

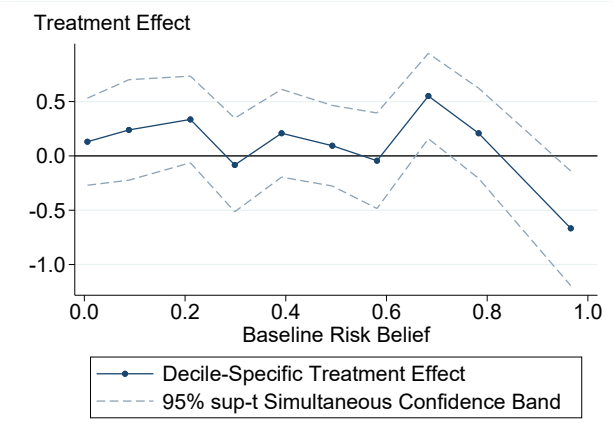

Notes: Sample is 1,292 people from 70 villages for whom both baseline and endline surveys were successfully completed. Baseline Risk Belief is the perceived chance of contracting HIV from a single unprotected sex act with a randomly chosen attractive person of the opposite sex from the local area, i.e. (Baseline Transmission Rate Belief) $\times$ (Baseline Prevalence Belief). The y-axis plots the total treatment effect for each decile of baseline risk beliefs, because Equation 1 has no omitted category of baseline beliefs and no treatment indicator. The x-axis shows the mean value of baseline risk beliefs for each decile.