Journal of Economic Behavior & Organization 118 (2015) 184–198

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

Risking it all for love? Resetting beliefs about HIV risk among low-income South African teens

Saugato Datta[a,*], Justine Burns[b], Brendan Maughan-Brown[c], Matthew Darling[a], Katherine Eyal[b]

[a] ideas42, 80 Broad Street, Floor 30, New York, NY 10002, USA
[b] School of Economics, University of Cape Town, Private Bag, Rondebosch 7701, Cape Town, South Africa
[c] Southern Africa Labour and Development Research Unit (SALDRU), University of Cape Town, Private Bag, Rondebosch 7701, Cape Town, South Africa

ARTICLE INFO

Article history:
Received 27 February 2014
Received in revised form 29 October 2014
Accepted 22 February 2015
Available online 3 March 2015

Keywords:
HIV
Risk perceptions
Age-disparate
Gender
Risk
Experiment

ABSTRACT

Research suggests that the much higher HIV prevalence among young women in sub-Saharan Africa than among males of their age cohort is linked to the high prevalence of age-disparate sexual partnerships, and that incorrect beliefs about the relationship between age and HIV-risk are partly responsible. We report the results of an experiment that tests whether a simple, computer-based "HIV risk game" leads to better understanding of the relationship between HIV-risk and age among low-income South African adolescents than a version of the traditional "brochure approach" to dispensing information does. Our results are striking. The randomly assigned treatment group, which receives repeated doses of information about the link between age and HIV-risk as feedback to their own responses to simple questions about relative HIV-risk, is significantly more likely to correctly identify which of a pair of hypothetical men or women of different ages is more likely to have HIV than the control group. Subjects in the treatment group answer, on average, 1.65 times as many questions about HIV risk and age correctly as those in the control group. We also find that subjects' (particularly female subjects') beliefs about HIV risk among women are less accurate than their beliefs about HIV risk among men. Finally, a follow-up survey with no significant difference in attrition rates between those in the treatment and control groups, shows substantially higher information retention among treatment subjects than among control subjects.

1. Introduction

Sub-Saharan Africa is at the center of the global HIV/AIDS epidemic. UNAIDS estimates that as of the end of 2012, approximately 35.3 million people worldwide were living with HIV; of these, approximately 25 million[1] lived in sub-Saharan Africa. Sub-Saharan Africa is thus home to about 71% of those living with HIV worldwide. Within Sub-Saharan Africa, South Africa has the largest number of people – an estimated 5.6 million – living with HIV in the region (UNAIDS, 2012a,b, UNAIDS, 2013a.

* Corresponding author. Tel.: +1 617 417 8626.
E-mail addresses: saugato@ideas42.org (S. Datta), justine.burns@uct.ac.za (J. Burns), brendan.maughanbrown@gmail.com (B. Maughan-Brown), matthew@ideas42.org (M. Darling), katherine.eyal@uct.ac.za (K. Eyal).

[1] Estimates for global prevalence range between 23.5 million and 26.6 million.

http://dx.doi.org/10.1016/j.jebo.2015.02.020
0167-2681/

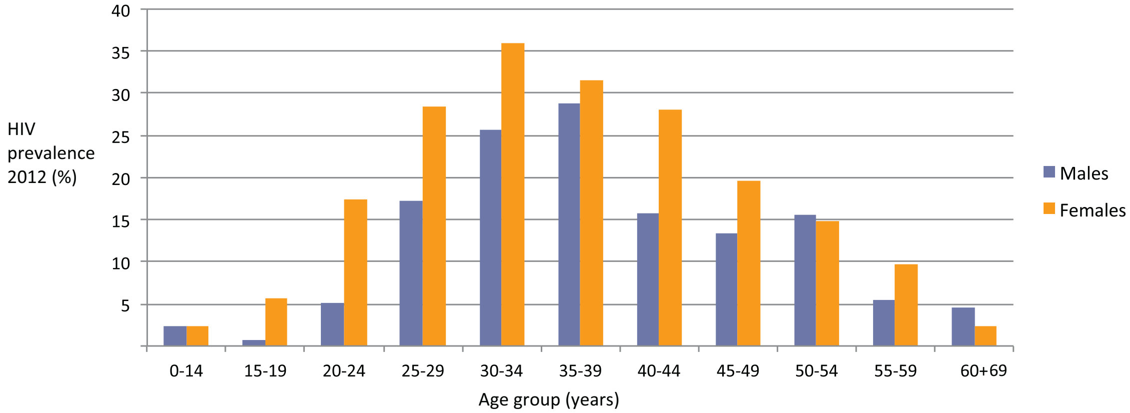

Fig. 1. HIV prevalence in South Africa by age and sex.

This overall pattern of high HIV prevalence in Sub-Saharan Africa and South Africa conceals large variations by gender and age. First, a disproportionate share of the burden of HIV falls on women, who account for 57% of all people living with HIV in sub-Saharan Africa (UNAIDS, 2013b). Secondly, the gender disparities in HIV prevalence in sub-Saharan Africa are not uniform with respect to age, but rather are *largest among younger cohorts*. HIV-infection risk is particularly acute for girls and young women compared with boys and young men of similar ages. In South Africa, for example, several studies have found that young women have HIV prevalence rates that range between two and four times those of men in their age cohort (Shisana et al., 2005, 2009; Nattrass et al., 2012; Pettifor et al., 2005).

While there is evidence that younger women are more physiologically vulnerable to HIV (see Royce et al., 1997; Varghese et al., 2002), the literature points to several non-physiological factors that help account for these stark gender inequalities in HIV prevalence among young adults. These include inequalities in education and economic opportunities, vulnerability to intimate partner violence, and women having sex with older men (UNAIDS, 2013b). In particular, a number of studies have identified age-disparate partnerships (defined in the literature as heterosexual partnerships in which there is a difference in age of five years or more between partners) as an HIV-risk factor (Katz and Low-Beer, 2008; Kelly et al., 2003; Gregson et al., 2002; Garnett and Anderson, 1993). Such relationships are widely prevalent across sub-Saharan Africa, as multiple qualitative and quantitative studies have shown (Luke, 2003). Two national surveys with data on age-mixing among South African women, for example, have shown that 32.6% of 15-to-24-year old women (Pettifor et al., 2005) and 36% of 16-to-24-year old women (Maughan-Brown et al., 2014) reported partnerships with an age-gap of at least five years. In a study among 21–45 year old men in Kenya, 70% of the sample reported an age-gap of five or more years with one of their recent non-marital partners, and 20% reported an age-gap of 10 or more years (Luke, 2005).

A substantial literature documents the relationship between HIV risk and age-disparate relationships. For example, Gregson et al. (2002) find strong empirical evidence that age-disparate partnerships contributed to the gender inequalities in HIV prevalence in rural Zimbabwe. Several South African studies corroborate this. For example, in the 2005 South African national HIV survey, HIV prevalence was 29.5% among 15–19 year old girls who had partners five years or more older than them, compared to the average HIV prevalence of 9.4% among girls in that age group (Shisana et al., 2005). More recently, young women in 2011 from the South African antenatal population whose partners were aged 15–19 had an average prevalence of 23.4% compared to 40.2% for those with partners aged 35–39 (South African Department of Health, 2012). While there are also studies (e.g., Harling et al., 2014 for rural KwaZulu-Natal) that fail to find such a link, the bulk of the available evidence indicates that age-disparate partnerships contribute to HIV infection risk.

Two mechanisms through which age-disparate partnerships increase HIV risk for women have been identified. First, the HIV-age profiles differ substantially among men and women, with HIV prevalence peaking five years later in men than among women (Shisana et al., 2009; Bärnighausen et al., 2008). In South Africa in 2008, for example, HIV-prevalence was 5.1% and 21% among 20-to-24-year old men and women respectively, and 15.7% and 33% among 25-to-29-year old men and women respectively (Shisana et al., 2009). Fig. 1 depicts HIV prevalence in South Africa in 2012 by age and sex. These differences in the relationship between age and HIV prevalence mean that women in age-disparate partnerships are, by definition, more likely to have sex with an HIV positive man than women who have partners of similar age. Second, there is evidence that partner age-gaps affect the behavior of partners within a relationship, particularly when it comes to behaviors related to safe sex. For example, partner age gaps have been found to reduce the likelihood that the younger partner will negotiate sex with a condom (Bankole et al., 2007; Longfield et al., 2004; Glynn et al., 2001; Luke, 2005; Langeni, 2007). These differences in turn make it more likely that HIV will be transmitted when there is a large age gap between partners than when there is no such gap.

Broadly speaking, the literature views age-disparate partnerships as being chosen when potential participants view them as being, on net, beneficial, in keeping with the neoclassical view of partner selection as laid out in Becker (1974), which models the choice of marriage partner as the result of a rational optimization exercise. However, the sources of this net positive benefit can vary. Some studies emphasize economic benefits. Thus, age-disparate relationships can be explained by noting that males traditionally earn more as they get older, so that females will often prefer older men because they can provide for them. Such considerations may be particularly important in contexts of poverty, where women and girls may select older partners to obtain household necessities and money for school fees (Buseh et al., 2002; Gregson et al., 2002). In other contexts, the economic benefits from age-disparate relationships involve the acquisition of luxury material goods and achievement of goals of social mobility, rather than simply meeting subsistence needs (Leclerc-Madlala, 2003).

A somewhat broader view expands the notion of benefits to include or emphasize those that are primarily psychological or cultural. A higher socioeconomic status may be viewed by women as status enhancing and be accompanied by a boost in self-esteem and self-confidence (Poulin, 2007). Older men may also be perceived to be better at meeting sexual and emotional (e.g., love, affection and affirmation) needs (Leclerc-Madlala, 2008). Culturally, given ties between fertility and masculine identities, age-disparate relationships can also be perceived by men to enhance social status and self-esteem. Furthermore, across Southern Africa, girls are still often encouraged to seek older partners as husbands as age-disparate marriages are perceived to be more stable (Leclerc-Madlala, 2008).

Finally, some researchers focus on risk perceptions, pointing out that whether individuals view age-disparate partnerships as beneficial on net depends on their perceptions of the risks associated with these partnerships (Leclerc-Madlala, 2008). Young women in South Africa have been found to perceive older men as more responsible and risk averse than younger men and therefore less likely to be living with HIV (Leclerc-Madlala, 2003). Similarly, for young women in Cape Town, Beauclair and Delva (2013) find that "a plurality . . . thought that dating an older man does not bring any adverse consequences". This increases the perceived attractiveness of such partnerships. Note, however, that these perceptions are inaccurate: among South African males aged 15–60, 15–19 year olds have the lowest HIV prevalence, followed by 20–24 year olds (Shisana et al., 2009). This is in keeping with a large literature on the existence of inaccurate risk perceptions (see Kahneman and Tversky, 1979), as well as evidence that individuals' perceptions of their own risk of contracting HIV can be remarkably inaccurate,[2] being disproportionately driven by recent, salient discussion of deaths due to HIV (see Sunstein, 2005).

Thinking about risk perceptions in the context of strong asymmetries in the HIV risk from age-disparate relationships vis-à-vis age-proximate ones raises interesting questions about whether individuals are informed about HIV risk and how they form and update their beliefs about the HIV risk associated with various kinds of partnerships (see Sunstein, 2005). All this depends, in turn, on where people receive information about HIV risk and whether this information is accurate, and the extent to which they update their beliefs to incorporate new information. Understanding how individuals form and update risk perceptions is thus potentially helpful in understanding age-disparate sex and thinking about ways to reduce it.

It is here – i.e., with respect to the formation and updating of risk perceptions about age-disparate sex – that this paper aims to make a contribution. We report the results from an experiment run with a sample of low-income adolescents in the Cape Town area. The purpose of our experiment was to test the effectiveness of an information-based intervention at correcting the prior beliefs of South African teenagers about the relative risks of being exposed to HIV depending on the age of their sexual partner. Subjects were randomly assigned to a treatment or control group. Control subjects read through a brief essay about HIV and sexual risk, which included a brief discussion of relative risks by age. In contrast, treatment subjects played ten rounds of a computer-based "HIV risk game". In each round, subjects were presented with the age and sex of two randomly generated individuals, and asked to choose which of the two was more likely to have HIV. Treatment subjects also received immediate feedback as to whether or not they had guessed correctly.

Our central results are as follows. First, we find that treatment subjects are significantly more likely to correctly identify which of two individuals is more likely to have HIV than control subjects are. The effect size is largest when female subjects answer the question about which of two women are more likely to have HIV. Secondly, we find that treatment subjects answer more questions about HIV-risk and age correctly than control subjects do. Only 7% of those in treatment get both questions asked wrong, compared with 35% of those in control; meanwhile, 63% of those in treatment get both questions right, compared with only 28% of controls. Regression estimates indicate that the treatment increases the mean number of questions a subject answers correctly by around 0.6.[3] Based on these results, we argue that playing a short "HIV risk game" with repetition and instant feedback leads to substantially more accurate perceptions of the relationship between HIV risk and age among our subject pool than equivalent information being provided through a more traditional "brochure approach", especially when initial beliefs are relatively less accurate.

The remainder of this paper is organized as follows. In Section 2, we review the existing evidence on interventions to tackle age-disparate sex and related behaviors among teens. Section 3 delves into the behavioral economics literature that helps us understand why teens might have incorrect priors about the HIV-risk posed by potential sexual partners of different ages, and briefly describes our findings from a series of focus groups carried out with low-income teens in Cape Town as a precursor to the experiment. Section 4 describes and provides the rationale for the intervention we ran. Section 5 describes how our experimental subjects were recruited and how the experiment was conducted. Section 6 describes our data and

[2] For example, using the Cape Area Panel Study, Anderson et al. (2007) find that the vast majority of young Black Africans perceived themselves to be at low or no risk of contracting HIV in the future. Given the extremely high prevalence of HIV among Black African women, it is clear that many of those surveyed in CAPS were severely underestimating their risk of getting HIV.

[3] Subjects in Treatment answer on average 1.55 questions correctly, compared with 0.94 among those in Control.

the outcome variables we use. Section 7 describes our key results. Section 8 concludes with a discussion of these results and implications for further research.

2. Evidence on reducing age-disparate partnerships

As outlined above, there are a number of different possible reasons why individuals may prefer age-disparate relationships. Each points toward different classes of interventions. The emphasis on economic benefits of age-disparate partnerships points toward the use of cash transfers, both conditional and unconditional, as a way of reducing the attractiveness of these partnerships. The role of social status and cultural capital point toward the need for communications and other interventions designed to reduce the social and cultural attractiveness of such partnerships. And finally, the emphasis on risk perceptions suggest the need for informational campaigns as well as the need for interventions that improve understanding of the HIV-risks associated with age-disparate partnerships. In keeping with this, researchers have used a variety of approaches to tackle the acceptability, desirability and prevalence of age-disparate sex and partnerships. While it is beyond the scope of this paper to be exhaustive, we summarize below a few notable studies that use one or other of these approaches.

To the extent that economic benefits may be an important driver of age-disparate partnerships, providing girls or young women with alternative sources of income or cash may reduce the benefits of such partnerships. Cluver et al. (2013) find that adolescent girls in South Africa who received a cash transfer under the government's child support grant were less likely to engage in transactional sex as well as age-disparate sex, although other risk behaviors were unaffected. In addition, a number of studies find that cash transfers, both conditional and unconditional, affect sexual behavior among adolescent girls (see Baird et al., 2009; Kohler and Thornton, 2012; Glennerster and Takavarasha, 2010; de Walque et al., 2012). Of particular interest is Baird et al. (2009), which investigated the effect of unconditional and conditional cash transfers to teenaged girls on school attendance and other outcomes, including marriage and fertility decisions using a randomized controlled trial. While not directly about age-disparate sex, the study found that teenage pregnancy and early marriage were both significantly reduced by the provision of unconditional transfers, implying that girls leaving school often do so for economic reasons, which also drive their decisions around marriage and sex.

Several interventions have sought to directly target the perceived social benefits of age-disparate sex through communications campaigns of various kinds, though few have been evaluated rigorously. A prominent example that has been carefully evaluated is the Fataki Campaign in Tanzania, a mass media campaign that alerted girls to the dangers of age-disparate relationships. A rigorous (albeit non-experimental) evaluation of the campaign finds strong evidence of exposure to the campaign having led to a lower likelihood of women engaging in cross-generational relationships (Kaufman et al., 2013).

Finally, several studies have focused on information provision about HIV. Here, the evidence is mixed. On the one hand, Sharp and Dellis (2010) argue that teens tend to display more than one risky behavior simultaneously and generally are fully aware of the consequences of these behaviors, explaining the frequent failure of interventions that are entirely knowledge-based. An additional concern, as Sharp and Dellis (2010) note, is that the school-based interventions they study often lead to an improvement in knowledge and attitudes but no accompanying change in behavior. On the other hand, Dupas (2011), who investigates the effect of informing teenagers in Kenya about the relative risks of sex with older men, finds dramatic effects on both proximate and long-term outcomes. The program resulted in a 65% decrease in pregnancy rates for teenaged girls involved with older men in the treatment group compared with the control group. Condom use also increased, and teenaged girls in the treatment group reported younger sexual partners than previously. One possible interpretation of the difference between this and other information-provision interventions is that general information or education is less effective at changing beliefs and behaviors than information that directly tries to correct an existing misperception that is driving behavior.

3. Changing beliefs about HIV risk and age

It is perhaps not that surprising that individuals may misperceive the risks associated with age-disparate relationships. As a starting point, note that estimating the likelihood that a person of a certain age has HIV necessarily involves estimating a probability. A large literature documents that individuals are in general remarkably poor at such tasks (see Kahneman, 2011 for a survey). In particular, however, Tversky and Kahneman (1973) posit that in cases where it is especially difficult to estimate a probability, individuals use a number of heuristics, or rules of thumb, including what is known in the literature as the "availability heuristic". Intuitively, the availability heuristic over-weights the probability of events that "come to mind first".

The availability heuristic plays a large role in how we think about risk – including sexual risks and in particular HIV risk (Sunstein, 2005). For example, studies in Kenya and Malawi have found that availability plays an important role in determining people's perceptions of the risk of contracting HIV. Risk perception is the product of discussions that "are often provoked by observing or hearing about an illness or a death" (Sunstein, 2005). In the context of our study, adolescents may be over-weighting the current, observable "promiscuity" of young individuals (which is visible and therefore available/"top of mind") while under-weighting the past sexual behavior of older individuals (which is not visible and therefore not similarly available). This may lead them to form estimates of relative risk which are based overwhelmingly on current behavior – which may indeed be more risky for younger men – while ignoring, or at least drastically under-weighting, the past sexual behavior of older men. The reason this leads to flawed estimates is of course that HIV risk is the result of lifetime behavior and exposure, not just current behavior.

Indeed, in a series of focus groups[4] run with adolescents from low income communities in Cape Town as a precursor to our experimental work, we found that it was common for participants to report incorrect beliefs about the relationship between HIV and age, beliefs that could plausibly lead young people to choose age-disparate partnerships, with possible repercussions for their own risk of acquiring HIV. Few adolescent girls or boys understood how HIV prevalence varied with age. Broadly speaking, the adolescents we interviewed had precisely the opposite understanding of the relationship between HIV risk and age than that revealed by the data on prevalence rates. This finding is consistent with other recent qualitative research from the Cape Town area (see Beauclair and Delva, 2013).[5]

In line with other studies (see Leclerc-Madlala, 2008), we therefore drew the conclusion that developing *effective* interventions that rapidly and effectively improve young women and older men's understanding of the HIV-risks associated with age-disparate partnerships is a priority.

4. Experimental design

The purpose of our experiment, programmed and conducted with the software z-Tree (Fischbacher, 2007), was to test the effectiveness of a novel information-based intervention at correcting the prior beliefs of South African teenagers about the relative risks of being exposed to HIV depending on the age of their sexual partner. Subjects were randomized into treatment and control groups. Control subjects read through a brief essay about HIV and sexual risk, which included a brief discussion of relative risks by age. As such, this approach mimics the more traditional information campaigns that rely on pamphlets or brochures to disseminate information.[6]

In contrast, treatment subjects played ten rounds of a computer-based "HIV risk game". In each round, subjects were presented with the age and sex of two randomly generated individuals, and asked to choose which of the two was more likely to have HIV. The age of the hypothetical individuals was uniformly distributed between 15 and 40. After they answered, subjects received immediate feedback as to whether or not they had guessed correctly, and were given the estimated prevalence for both individuals, based on data from the South African National HIV, Behaviour and Health Survey 2012. If both individuals were male, they also received a hint directly alerting them to the fact that older men were more likely to have HIV.[7] Rather than being provided information directly (as our control group was), our treatment group received the information about HIV-risk and age as part of feedback to their responses to questions about HIV-risk and age.

The fundamental justification for the approach embodied in our intervention comes from the literature on learning and updating of priors. We assume that subjects enter the experiment with some priors about HIV prevalence rates in the population, and how these vary by gender. Once subjects are exposed to information about actual HIV prevalence rates by gender in the population, it seems reasonable that they would incorporate this new evidence to update their prior accordingly, albeit imperfectly, subject to some sort of psychological biases.

Note that we do not make any specific claims about the manner in which subjects might learn in this setting. Indeed, the existing literature suggests that subjects are sufficiently heterogeneous so as to render a single theory of decision-making or learning inadequate (El-Gamal and Grether, 1995), and there is substantial evidence to suggest that individuals are imperfect Bayesian updaters at best (Bartoli, 2014). Using a general classification procedure to describe the most likely collection of rules used by subjects in experimental settings, El-Gamal and Grether (1995) argue that subjects tend to rely on Bayesian inference in combination with a representativeness heuristic and conservatism heuristic.[8]

Of course, both subjects in treatment and control groups receive information about HIV prevalence by age and therefore have an opportunity to update their priors. However, the literature provides several reasons to think that the treatment would be more effective at aiding subjects in updating their priors than the straightforward information provision via the "brochure approach" received by control group subjects.

First, the control group only receives a single "dose" of HIV-related information, whereas the treatment group answers a series of questions about HIV risk and age sequentially. Second, the control group receives no feedback, whereas the treatment group receives immediate feedback as to whether they have answered correctly or not.

---

[4] These focus groups delved into subjects' attitudes toward sexual relationships (age-disparate and otherwise), their beliefs about HIV risk, and the drivers of their sexual behavior.

[5] Other misperceptions about HIV transmission undoubtedly also exist. For example, a recent paper (Maughan-Brown et al., 2014) shows that men and women in Malawi vastly overestimate the per-sexual-encounter probability of HIV infection. However, we are not aware of any evidence that such incorrect beliefs vary by the age of the sexual partner, and so we do not focus on them here. Further, incorrect beliefs about the relationship between HIV prevalence and age were the most salient in our focus groups as well as in the literature.

[6] Note that our experiment does not have a "true control group", in that even our control group does read the brochure. The implications of this for our interpretation of control groups' beliefs, and an alternative measure of baseline beliefs, are discussed in Footnote 21 and also in the Supplementary Material.

[7] This was done only for male–male pairs because the primary goal was to alter understanding of the risk of sex with older men, as well as to enable us to check whether subjects displayed better learning about risk among men (for whom they received a direct hint) or whether this hint made no difference. The program was hard-coded to ensure that every subject received at least two male-male pairs.

[8] Similarly, Barberis et al. (1998) and Griffin and Tversky (1992) present evidence in favor of these two heuristics in individual decision making.

These two features of the intervention are important in that they can plausibly be thought to affect the likelihood of priors being updated. There are several related reasons for this. First, the feedback received by the treatment group allows for learning to be reinforced or errors to be corrected. This arguably reduces the cost of learning relative to that experienced by the control group. This in turn means that updating is more likely for the treatment group than for the control group.

Secondly, many decision-making biases which have been observed in laboratory settings have not been found to be applicable in field research with experienced market participants. For example, while the endowment effect (i.e., the hypothesis that people ascribe more value to things merely because they own them) has been observed in numerous laboratory settings (Kahneman et al., 1991) the effect is substantially reduced – sometimes eliminated – when the subjects are experienced traders (List, 2003). This suggests that repeated decision-making is likely to attenuate the problem of biased priors. However, such repeated decision-making is not always possible, particularly for decisions that are made only rarely (such as purchasing a house) or where experimentation is risky (as with sexual risk). In such a setting, repeatedly answering questions about the relationship between HIV risk and age provides a way to simulate repeated decision-making, albeit virtually, and should cause biases in priors to attenuate.

Finally, any training on HIV risk would need to be retained over a long time period in order to affect long run-behavior. A broad literature (see Bertsch et al., 2007 for a review) shows a consistent "generation effect" occurs when laboratory subjects generate new information rather than read information passively. This provides another reason to be optimistic that those in our treatment group would learn and retain more information about HIV risk and its relationship with age than those in our control group.

5. Subject recruitment and experimental protocol

The experiment was advertised using posters at and around a public library[9] in Khayelitsha, Cape Town. Khayelitsha is a densely populated semi-informal township in Cape Town, South Africa, of nearly 400,000 people. The library is in a high traffic area, and is frequented by large numbers of teenagers from the area, which is predominantly (98.6%) Black African. The recruitment literature and the research assistants who signed potential participants up emphasized that the experiment was targeted at 15–19 year olds.[10] After recruitment, subjects went through the informed consent process. Subjects were provided with consent forms in both English and Xhosa, and a hard copy was given so that subject could inform their parents about the study. The consent materials (as well as recruitment materials) informed potential subjects that they would receive ZAR50 (approximately US$5) for their participation. Whilst the sessions were not timed, they lasted 70 min on average, with the bulk of time spent on set-up and the post-intervention survey.

The experiment took place over four days, with a total of 9 sessions. For each session, subjects were brought into a room with a number of laptop computers onto which the experiment had been loaded. Each participant sat in front of a laptop computer. Participants were randomly assigned to either treatment or control status by z-Tree. The (simple) instructions[11] were built into the program itself. In addition, the research assistant conducting the session explained the instructions to participants in both Xhosa and English, and these assistants were available throughout the sessions to answer questions of clarification. Students first completed a short questionnaire, which asked them for their demographic details, and then proceeded onto the experimental portion. In that portion, control subjects read through a brief essay about HIV and sexual risk, which included a brief discussion of relative risks by age. Treatment subjects played 10 rounds of the "HIV risk game" described in Section 4.

Once both treatment and control subjects had completed their relevant tasks, both groups completed a short survey, which asked them the questions about HIV risk and age which are the basis of our key outcome variables.[12] These questions were: Is a 20-year old man (woman) or a 30-year-old man (woman) more likely to have HIV? All subjects answered both questions (i.e., one question comparing two men, and another comparing two women).[13] This concluded the experiment.

Furthermore, approximately three months after participants had participated in the study, they were recontacted by phone, and once again asked to say whether a 20-year-old man or a 30-year-old man were more likely to have HIV. This short follow-up interview was undertaken in order to assess whether or not the treatment intervention had any persistent effect on understanding of HIV risk. There were no monetary incentives for correct answers in this follow-up survey.

[9] Our sample, which was drawn from among teenagers who frequented a public library, may have been better-informed and more socially connected than the average teenager.

[10] In a few cases, research assistants turned away subjects who indicated that they were well outside the target age range.

[11] Controls read: "Please read the following short description of HIV. There will be a short test at the end! Treatments read:" We are going to ask you several short questions about who is more likely to have HIV. Please make your best guess!".

[12] Subjects also answered questions about the prevalence of and attitudes toward age-disparate sex in their community, their preferences about the ideal age of their partner, and open-ended questions about what they took away from the experiment.

[13] Our decision to ask only a few questions to evaluate subjects' understanding of the relationship between HIV risk and age was based on the following: (1) we wanted to present choices that represented some approximation of age-proximate and age-disparate partners (but below peak age) for 15–20 year olds, the age of our subjects and broader target population. Taking into consideration the age of our target population, and the fact that HIV prevalence peaks at 30–34 for men and 25–29 for women, this suggested that limiting the age comparisons of the hypothetical individuals in the post-survey to lie between 20 and 30 years of age was reasonable. (2) Limiting the comparisons to same gender pairs allows us to identify the extent to which subjects are able to consider the relative risk of engaging in a relationship with members of a specific gender as age varies. (3) Since subjects were also asked to answer a number of other survey questions on their demographic and socio-economic background, we wanted to avoid subject fatigue.

Table 1
Summary statistics, experimental sample.

All subjects	
Variable	Value
Age (n = 156)	
Mean age	17.4
Median age	18
Race (n = 162)	
% Black African	96.9
% Colored	1.9
% Indian	1.2
% White	0
% Mixed race	0
Gender (n = 162)	
% Male	58.6
% Female	40.1
% Other	1.2

Table 2
Characteristics of individuals in treatment and control groups.

	Control	Treatment	Difference (T – C)
Age	17.44	17.43	−0.01
	(0.16)	(0.17)	(0.23)
n	79	77	156
% Male	0.60	0.57	−0.03
	(0.05)	(0.05)	(0.08)
n	80	82	162
% Black	0.96	0.98	0.01
	(0.02)	(0.02)	(0.03)
n	80	82	162

Notes: (1) Numbers in parentheses are standard errors. (2) For age, $n = 79$ in treatment and 77 in control as 6 subjects did not enter an exact age.
* Statistical significance at $p < 0.10$.
** Statistical significance at $p < 0.05$.
*** Statistical significance at $p < 0.01$.

6. Characteristics of experimental sample and treatment and control groups

Our sample consists of 162 individuals recruited as described in Section 5. Table 1 summarizes the information collected on the experimental sample's age, ethnicity, and gender.[14] 151 of the 156 subjects who entered their age in years were within the target age range of 15–19 years, with 5 others being a year older the desired age range.[15] According to South Africa's 2011 national census, 98.6% of Khayelitsha residents were Black African, so the ethnic mix of the experimental sample (96.9% Black African) reflects the demographics of the site rather than any sampling bias.[16] The experiment also attracted substantially more male than female participants, with 58.6% males in the overall sample. Broadly speaking, we are satisfied that our sample pool was suitable for the experiment, in that it was dominated by teenagers in the target age range and its racial mix was similar to that of the catchment area of the experiment site.

6.1. Randomization check: no significant differences between treatment and control

Since subjects in each session were randomly assigned to either a treatment group or a control group, we should see no significant differences between the average values of the demographic characteristics summarized in Table 1 between our randomly assigned treatment and control groups as long as our randomization was not compromised.

Table 2 compares subjects who were randomly assigned to receive the treatment (i.e., play the "HIV risk game") to those who were randomly assigned to read the control text. As it shows, there were no significant differences in demographic

[14] As part of the survey built into the experiment, all participants had to enter their age, ethnicity, and gender. They also entered their grade level, which we do not report because it was highly correlated with age when reported, but was not available for 15 out-of-school subjects.
[15] We do not accurately know the ages of 6 individuals who chose either "older than 20" or "younger than 15".
[16] The categories we use to code ethnicity conform to official practice in South Africa. It is worth noting that the persistence of apartheid-era spatial segregation by race means that many townships are, like Khayelitsha, dominated almost entirely by members of one ethnicity.

Table 3
Percent of subjects answering questions on HIV risk-age relationship correctly.

	Control	Treatment	Difference (T – C)
Panel A: all subjects			
% Correctly identified older man as riskier	0.63	0.80	0.18**
	(0.05)	(0.04)	(0.07)
% Correctly identified older woman as riskier	0.30	0.76	0.46**
	(0.05)	(0.05)	(0.07)
n	80	82	162
Panel B: male subjects			
% Correctly identified older man as riskier	0.52	0.77	0.25**
	(0.07)	(0.06)	(0.10)
% Correctly identified older woman as riskier	0.35	0.79	0.43***
	(0.07)	(0.06)	(0.09)
n	48	47	95
Panel C: female subjects			
% Correctly identified older man as riskier	0.77	0.85	0.08
	(0.08)	(0.06)	(0.10)
% Correctly identified older woman as riskier	0.23	0.71	0.48***
	(0.08)	(0.08)	(0.11)
n	31	34	65

Note: Numbers in parentheses are standard errors.
** Statistical significance at $p < 0.05$.
*** Statistical significance at $p < 0.01$.

characteristics – age, gender, and race/ethnicity – between those in the treatment and control groups.[17] The control group had slightly fewer Black Africans (96%) than the treatment group (98%), and slightly more males (60%) than the treatment group (57%), but neither of these differences are statistically significant at conventional levels.

Based on this, we are therefore confident that our randomization was valid. This means that any observed differences between the outcomes for treatment and control groups are attributable to the treatment, rather than to pre-existing differences between individuals in either group.

6.2. Outcome variables: measuring subjects' understanding of relationship between HIV-risk and age

Experimental subjects' answers to two questions asked in the post-intervention survey allow us to get at their understanding of the relationship between age and HIV risk, the focus of our study. As discussed in Section 5, the first of these questions asked them to identify which of a 20-year-old man and a 30-year-old man were more likely to have HIV. The second asked the identical question but changed the gender of the people the question was about. The pattern of HIV prevalence by age in South Africa means that the correct response in either case would be to identify the older man/woman as being more likely to have HIV.

We therefore construct three outcome variables that measure subjects' ability to correctly answer questions on HIV risk and age. The first is a binary variable that takes the value 1 if a subject correctly answered the HIV risk question about two men, and takes the value 0 otherwise. The second is a binary variable that captures whether the subject correctly answered the equivalent question about which of two *women* was more likely to have HIV. These variables are of independent interest, since subjects might have quite different priors about HIV prevalence and age among men and women. However, we were also interested in a measure of overall "correctness", since the underlying principle that a subject must grasp to answer either question is identical. Our third outcome variable is thus the *number of HIV-risk questions* – 0, 1 or 2 – that an individual answered correctly.

7. Results

7.1. Result 1: treatment subjects are significantly more likely to correctly identify an older person as being more likely to have HIV

Panel A of Table 3 shows that treatment subjects are significantly more likely to identify the older of a pair of individuals as more likely to have HIV. As the top row shows, 80% of subjects in the treatment group correctly identified the older of a pair of men as being more likely to have HIV. This is 18 percentage points higher than the corresponding fraction for the control group. This difference is both large in magnitude and statistically significant. The results are qualitatively similar,

[17] There is also no significant difference in the school grade level for the 147 subjects who were in school. The mean for treatment, conditional on being in school, was 11.47, whereas that for control was 11.52.

Table 4
Effect of treatment on probability of correctly identifying older individual as riskier.

Independent variable	I	II	III
Panel A: probability of correctly identifying older man as riskier			
Treatment	0.18**	0.16**	0.19**
	(0.07)	(0.07)	(0.09)
Treatment × Female			−0.07
			(0.14)
Age	N	Y	Y
Gender	N	Y	Y
Ethnicity	N	Y	Y
Session	N	Y	Y
n	162	162	162
Panel B: probability of correctly identifying older woman as riskier			
Treatment	0.46***	0.44***	0.38***
	(0.07)	(0.07)	(0.09)
Treatment × Female			0.14
			(0.14)
Age	N	Y	Y
Gender	N	Y	Y
Ethnicity	N	Y	Y
Session	N	Y	Y
n	162	162	162

Notes: (1) The dependent variable is an indicator variable which takes the value 1 if the subject correctly identified an older man/woman as having a higher probability of being HIV-positive. (2) All regressions are OLS, with robust standard errors clustered at the level of the individual experimental subject.
** Statistical significance at $p < 0.05$.
*** Statistical significance at $p < 0.01$.

but larger in magnitude, when we turn to the analogous question asking subjects to identify which of a pair of women of different ages was more likely to have HIV. In this case, 76% of those in the treatment group answered correctly, while only 30% of those in the control group did. Once again, the difference is large, at 45 percentage points, and statistically significant.

Panels B and C of Table 3 present these results separately by the gender of the experimental subject. We find that while there are no *qualitative* differences by gender, there are indeed differences in the magnitude of the difference between treatment and control depending on the gender of the participant. Males in the treatment group are significantly more likely to identify both an older man and an older woman as being more likely to have HIV than males in the control group are. However, while females in the treatment group are significantly more likely to identify an older woman as more likely to have HIV than females in the control group, the corresponding difference in the case of the comparison between two men is not statistically significant. We discuss the reasons for this and what they imply for our overall findings as part of the discussion of Result 3 below.

Table 4 presents the regression counterparts of these results. Panel A shows the results for regressions where the dependent variable is the probability of subjects' correctly identifying the older of a pair of men as more likely to have HIV. Panel B shows the equivalent results when the dependent variable measures the probability of correctly identifying the older woman in a pair as more likely to have HIV. In all cases, the coefficient of interest is that on the treatment status. All regressions are OLS, with robust standard errors clustered at the subject level.[18] The basic specification (Column I of each panel) has no additional controls; Column II of each panel adds controls for demographic variables (age, gender, and race); and Column III of each panel adds an interaction between gender and treatment status, to control for the differential performance of male and female subjects discussed earlier.

Reading across the top row of Table 4, we see that the coefficient of interest is positive and statistically significant in all specifications, implying that the treatment significantly increases the probability of subjects answering the questions about HIV risk and age correctly. The effect size is larger for the HIV-risk question for women of different ages than for men of different ages. In the former case (Panel B) it varies between 0.44 and 0.46, while in the latter case (Panel A) it ranges between 0.16 and 0.19. Within each panel, effect size is consistent across specifications, with no loss of statistical significance.[19]

7.2. Result 2: treatment group answers more questions correctly than control group

Table 5 shows the distribution of the number of correct answers to the HIV-risk questions for members of the treatment and control groups. Panel A shows the distribution for all subjects, Panel B for male subjects, and Panel C for female subjects. As we see from Panel A, those in the treatment group were 28 percentage points less likely to get no answers right than

[18] We also ran probit regressions, which are not reported here. There was no qualitative difference in the result: the coefficient on the treatment variable was always positive and significant.

[19] It is worth noting that the treatment effect is larger when subjects are asked about women. However, recall that subjects received a hint when faced with a male–male pairing to the effect that "Older men are riskier" but did not receive any hint in relation to female pairings. This suggests that the direct hint was not solely what drove learning in the experiment, but rather the nature of the treatment itself.

Table 5
Number of correct responses to HIV-age questions.

# Correct	Control	Treatment	Difference (T – C)
Panel A: all subjects			
0	0.35	0.07	−0.28***
	(0.05)	(0.03)	(0.06)
1	0.38	0.29	−0.08
	(0.05)	(0.05)	(0.07)
2	0.28	0.63	0.36***
	(0.05)	(0.05)	(0.07)
n	80	82	162
Panel B: male subjects			
0	0.44	0.09	−0.35***
	(0.07)	(0.04)	(0.08)
1	0.25	0.28	0.03
	(0.06)	(0.07)	(0.09)
2	0.31	0.64	0.33***
	(0.07)	(0.07)	(0.10)
n	48	47	95
Panel C: female subjects			
0	0.23	0.06	0.17*
	(0.08)	(0.04)	(0.09)
1	0.55	0.32	−0.22*
	(0.09)	(0.08)	(0.12)
2	0.23	0.62	0.39***
	(0.08)	(0.08)	(0.11)
n	31	34	65

Note: Numbers in parentheses are standard errors.
* Statistical significance at $p < 0.10$.
*** Statistical significance at $p < 0.01$.

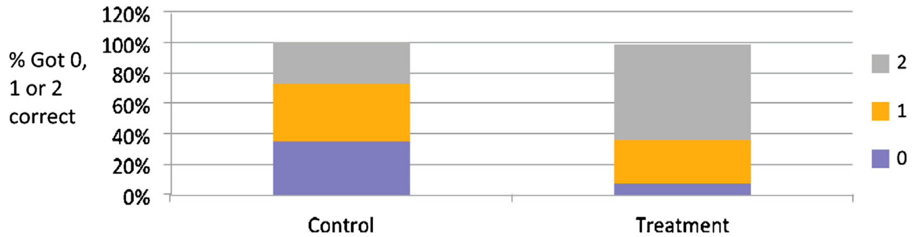

Fig. 2. Distribution of number of correct answers, control versus treatment.

those in the control group. They were also 36 percentage points more likely to get both answers right. Fig. 2 shows the distribution of the number of correct answers for treatment and control groups side by side. Fig. 3 provides an alternative visual representation, plotting the cumulative percentage of treatment and control subjects by the number of incorrect answers. Both serve to visually illustrate the shift in the distribution of the number of (in)correct answers.

This pattern persists when we look at male subjects (Panel B) and female subjects (Panel C) separately. In all cases, those in the treatment group are significantly more likely to answer both HIV-risk questions correctly and significantly less likely

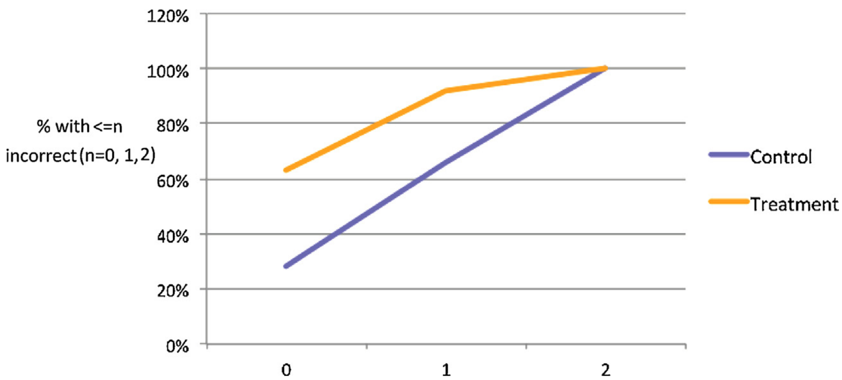

Fig. 3. Cumulative distribution of number of incorrect answers.

Table 6
Effect of treatment on number of correct answers to age-HIV risk questions.

Independent variable	I	II	III
Treatment	0.64***	0.60***	0.57***
	(0.11)	(0.11)	(0.15)
Treatment × Female			0.07
			(0.22)
Age	N	Y	Y
Gender	N	Y	Y
Ethnicity	N	Y	Y
Session	N	Y	Y
n	162	162	162

Notes: (1) The dependent variable is the number of age-HIV risk questions the subject answered correctly. (2) All regressions are OLS, with robust standard errors clustered at the level of the individual subject.
*** Statistical significance at $p < 0.01$.

to get both wrong. The difference between treatment and control groups in the fraction of those getting one answer right is generally not significant except in the case of female subjects, where it is negative and marginally significant.

Table 6 provides the regression counterpart of these results, where our dependent variable is the number of correct answers. As before, we run linear regressions with robust errors clustered at the subject level. In all specifications – with no controls (Column I), demographic controls (Column II) and demographic controls together with the gender-treatment interaction (Column III), the coefficient of interest, i.e., the coefficient on treatment status, is positive and statistically significant. The effect size is large: our estimates suggest that the treatment increases the mean number of correct answers by approximately 0.6. Given that controls get on average 0.94 questions right, this implies that the treatment increases the average number of correct answers by just under two-thirds.

7.3. *Result 3: control subjects' beliefs about women are less accurate than their beliefs about men, and this difference is more marked for female subjects*

Our experiment was not intended to delve into differences between the accuracy of subjects' beliefs about HIV risk among women and their beliefs about HIV risk among men. However, returning to Table 3, but looking now at the same set of subjects' responses to the question about men and the question about women, we see a striking difference in the accuracy of subjects' beliefs (proxied here by the responses of the control group) about HIV risk among men and among women, and that this difference is larger for female subjects than for male subjects.[20]

Note that 63% of controls correctly identified an older man as being more likely to have HIV, while only 30% of them correctly identified an older woman as being more likely to have HIV.[21] A test of proportions shows[22] that the difference between controls' probability of getting the "male question" right and their probability of getting the "female question" right is statistically significant at the 1% level. The control group thus has more inaccurate beliefs about HIV risk among women than they do about HIV risk among men. Alternatively, their beliefs about HIV risk among men are more accurate than their priors about HIV risk among women.

Crucially, this difference is even larger when we restrict ourselves to only looking at female control subjects. 77% of control group females identify the older man as more likely to have HIV, while only 23% correctly identify the older woman as more likely to have HIV. These proportions fail a test of equality at the 1% level. Females in the control group thus have strikingly more accurate beliefs about HIV-risk among women than they do about HIV-risk among men.

This finding suggests that there was far less scope for any intervention to "reset" the beliefs of female subjects about HIV-risk among men. It thus helps explain why we see no significant treatment effect on females' ability to identify the man more likely to have HIV.

[20] An obvious concern is that the measures of control group beliefs could reflect the result of learning from the brochure, given that we have no "true control group" that was exposed to neither the brochure nor the game. In the supplementary material, we discuss an alternative measure of baseline beliefs, based on the initial responses of the treatment group. Since our measure of "baseline beliefs" is very similar to the measure based on the control group's response to the post-brochure survey, we conclude that the brochure's effect was negligible and that we are justified in using our control group's responses to proxy for baseline beliefs.

[21] The only other measure of related baseline beliefs we have found in the literature is from Dupas (2011). There, 29% of girls and 25% of boys correctly believed that a 25-year old man was more likely to have HIV than a teenaged boy. While this is quite different from our control group's beliefs, as well as our proxy measures for baseline beliefs based on the treatment group's initial responses, we should note that the Dupas (2011) data is for a representative cohort in a much lower-prevalence setting than ours, and predates our data by 9 years. In addition, our sample, which was drawn from among teenagers who frequented a public library, may have been better informed and more socially connected than the average teenager.

[22] The z-scores (p-values) for a test of equality of proportions for All Controls/Male Controls/Female Controls are 4.18 (0.000), 1.68 (0.093), and 4.25 (0.000) respectively.

Table 7
Characteristics of individuals in follow-up sample.

	Control	Treatment	Difference (C − T)
Age	17.95	17.71	0.23
	(0.23)	(0.26)	(0.35)
n	38	28	66
% Male	0.55	0.56	−0.01
	(0.08)	(0.09)	(0.12)
n	38	32	70
School grade	11.17	11.19	−0.02
	(0.20)	(0.17)	(0.27)
n	30	27	57
% Black	1	0.97	0.03
	(0)	(0.03)	(0.03)
n	38	32	70

Note: Numbers in parentheses are standard errors.
* Statistical significance at $p < 0.10$.
** Statistical significance at $p < 0.05$.
*** Statistical significance at $p < 0.01$.

Table 8
Percent correct by treatment and control status, follow-up survey.

	Correctly identified older man as riskier		
	Control	Treatment	Difference
All subjects	0.61	0.78	0.18
	(0.08)	(0.07)	(0.11)
n	38	32	70
Just male subjects	0.71	0.72	0.01
	(0.1)	(0.11)	(0.14)
n	21	18	39
Just female subjects	0.47	0.86	0.39**
	(0.12)	(0.09)	(0.15)
n	17	14	31

Note: Numbers in parentheses are standard errors.
** Statistical significance at $p < 0.05$.

7.4. Result 4: treatment group subjects, especially women, remain more likely to identify older male as more risky in follow-up survey

As discussed in Section 5, experimental subjects provided their mobile numbers and consent to be contacted for a brief follow-up survey. Approximately three months after the original experiment, research assistants attempted to contact all participants using the numbers provided. They succeeded in contacting 70 out of 162 participants, a success rate of 43.2%. Table 7 shows that the "original treatment subjects" whom we succeeded in re-contacting look remarkably similar demographically to the "original control subjects" who responded to our follow-up survey. We see no statistically significant differences in age, race/ethnicity or gender. Nor do we see any significant correlation between attrition and treatment status. Indeed, while it is often the case that treatment subjects are less prone to attrition than control subjects, the opposite is true in our case: attrition is higher (at 61%) for our treatment group than it is for our control group (51%), though this difference is not significant at conventional levels.

Table 8 shows the results of our follow-up survey, during which respondents were once again asked to say whether a 20-year-old man or a 30-year-old man were more likely to have HIV. Three months after the intervention, those of our original treatment group who responded to the follow-up survey were 18 percentage points more likely to correctly identify an older man as being more likely to have HIV than those formerly in our control group. Interestingly, the largest effect is among female respondents, for whom our original treatment was weakest. 85% of females who had been in the treatment group answered the question correctly, compared with a mere 47% of former controls.

What should we make of these results? One possibility is that the treatment really did have persistent effects on understanding of HIV risk. Our finding that there is no differential attrition between treatment and control groups supports this possibility.[23]

[23] See Supplementary Material for regressions of attrition status on treatment status.

Table 9.1
Control subjects: performance in original experiment by inclusion in follow-up sample.

	Control not in follow-up	Control in follow-up	Difference (N – Y)
% Got man question right in lab	0.55	0.71	−0.16
	(0.08)	(0.07)	(0.11)
% Got woman question right in lab	0.26	0.34	−0.08
	(0.07)	(0.08)	(0.10)
# of questions correct	0.81	1.05	−0.24
	(0.12)	(0.12)	(0.18)
n	42	38	80

Note: Numbers in parentheses are standard errors.
* Statistical significance at $p < 0.10$.
** Statistical significance at $p < 0.05$.
*** Statistical significance at $p < 0.01$.

Table 9.2
Treatment subjects: performance in original experiment by inclusion in follow-up sample.

	Treatment not in follow-up	Treatment in follow-up	Difference (N – Y)
% Got man question right in lab	0.76	0.88	−0.12
	(0.06)	(0.06)	(0.08)
% Got woman question right in lab	0.72	0.81	−0.09
	(0.06)	(0.07)	(0.09)
# of questions correct	1.48	1.69	−0.21
	(0.09)	(0.10)	(0.14)
n	50	32	82

Note: Numbers in parentheses are standard errors.
* Statistical significance at $p < 0.10$.
** Statistical significance at $p < 0.05$.
*** Statistical significance at $p < 0.01$.

However, note that the results in Table 8 might still be driven by selection if our follow-up survey was disproportionately likely to include those in the treatment group who had particularly accurate understanding of HIV risk, while showing no (or a smaller) such bias in the case of control participants.

Tables 9.1 and 9.2 suggest that there is indeed some positive selection into our follow-up sample, and thus the need for a note of caution in interpreting the results in Table 8. In these tables, we compare the original answers to the HIV risk and age questions for "responsive" members of the control group to those who did not respond (Table 9.1), and do a similar analysis for responsive and non-responsive members of the treatment group (Table 9.2). We see that those we succeeded in contacting – from either group – had been more likely to correctly identify an older man and an older woman as more likely to have HIV than those who attrited. As expected, both treatment and control subjects in our follow-up sample had answered more "HIV-risk and age" questions correctly in the initial experiment than those who did not respond. None of these differences is significant at conventional levels, due in part to small sample sizes, but we note that the sign of the difference does indicate positive selection (i.e., those who had originally done better were more likely to be in our follow-up sample, but this does not differ for treatment versus control).

This suggests that the difference in recall is not being driven by differential baseline ability or attrition.

8. Discussion

We find that a very simple game, where subjects are asked to repeatedly answer questions about relative HIV-risk does substantially better at passing on information about the link between HIV and age than a traditional brochure approach.[24] While the game is simple, it provides subjects with repeated exposure to information and immediate feedback about their responses. In addition to the immediate effects on comprehension, which are large, we also find some evidence that those who played the HIV game also retained the information about the relationship between HIV-risk and age better: three months after the experiment, they were more likely to answer a question about HIV risk correctly than those who received the information "brochure".

Our results are in line with the literature on "generation effects", which suggests that information is more readily comprehended and retained when it is generated by individuals themselves, rather than provided ready-made. These results have implications for the way in which information is delivered. They suggest that "gamifying" information acquisition may

[24] We find no evidence that our effects can be attributed to a pure "dose effect", i.e. from treatment individuals spending more time on the treatment than controls did on the brochure. In fact, as we discuss in the Supplementary Material, controls 342 s on average on the brochure, more than twice the 133 s on average that treatments spent playing the HIV game.

be a far more effective way to deliver information about matters – such as HIV risk – which is important for individuals to understand but which many do not, with adverse consequences for their lives.

Of course, we cannot say whether our experiment affected more than knowledge. In particular, we cannot say whether having understood that age-disparate partnerships are riskier than age-proximate ones leads to changes in subjects' propensity to engage in age-disparate partnerships. However, past research has shown that increased understanding of relative HIV risk has dramatic effects on actual behavior. Dupas (2011), which is closely related to this paper in terms of the key hypothesized mechanism by which the intervention works,[25] finds that informing teenagers in Kenya about the relative risks of sex with older men has dramatic effects on both proximate and long-term outcomes. The program she evaluates resulted in a 65% decrease in pregnancy rates for teenaged girls involved with older men in the treatment group compared with the control group. Teenaged girls in the treatment group reported younger sexual partners than previously, but condom use also increased, so that this did not lead to more pregnancies from similar-age pairings while reducing those from age-disparate pairings.

One further concern merits our attention. An intervention such as this one in effect tries to shift the sexual activity of teenaged females away from older males toward males nearer their age. This could, in theory, raise HIV rates among younger males, especially if there is no compensating change in complementary behaviors, such as condom use. This is a legitimate concern, and one that deserves serious attention in future studies of intervention of this nature that measure behavioral and health outcomes. However, several results – the finding about increased condom use in similar-age pairings in Dupas (2011), the finding that partner age gaps are likely to reduce the likelihood that the younger partner will negotiate sex with a condom (Bankole et al., 2007; Longfield et al., 2004; Glynn et al., 2001; Luke, 2005), and the results of a large study among men in Botswana that found that there was a 28% increase in the odds of having unprotected sex for each one-year increase in the age-gap between partners (Langeni, 2007) – suggest that young women shifting toward younger partners are in fact more likely to use protection. This goes some way toward obviating the concern about the effects of our intervention on HIV rates among young men.

The prototype game we test here conveys much the same information as the classroom education in Dupas (2011) but is potentially scalable at very low cost, since it is a simple computer-based game that can be played with little or no supervision. Based on discussions with local software developers, we estimate that a version of this game that could be installed on existing computers in school computer laboratories could be developed and installed for within $10,000. Further, data from the Western Cape Government Education Department allows us to calculate the number of students in Grades 8–12 who could be reached using this method[26] as 141,000 for the entire province, or 83,500 for the Cape Town metropolitan area, giving us an estimated cost per student reached of between $0.07 and $0.12, much lower than the cost of $1 per student reached in Dupas (2011).[27]

The results from our experiment are encouraging for future attempts to develop a simple but scalable intervention that could cost-effectively leverage the benefits of gamification to substantially increase young people's understanding of the risks inherent in age-disparate partnerships. Further research should focus on measuring effects on longer-term learning and retention as well as behavioral change in the domain of risky sexual behavior, care-seeking, teen pregnancy and other related outcomes.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.jebo.2015.02.020.

References

Anderson, K.G., Beutel, A.M., Maughan-Brown, B., 2007. HIV risk perceptions and first sexual intercourse among youth in Cape Town, South Africa'. Int. Fam. Plan. Perspect., 98–105.

Baird, S., Chirwa, E., McIntosh, C., Ozler, B., 2009. The Short-term Impacts of a Schooling Conditional Cash Transfer Program on the Sexual Behavior of Young Women. World Bank Policy Research Working Paper 5089. World Bank, Washington, DC.

Bankole, A., Ahmed, F.H., Neema, S., Ouedrago, C., Konyani, S., 2007. Knowledge of correct condom use and consistency of use among adolescents in four countries in sub-Saharan Africa. Afr. J. Reprod. Health 11 (3), 97–220.

Barberis, N., Shleifer, A., Vishny, R., 1998. A model of investor sentiment. J. Financ. Econ. 49, 307–343.

Bärnighausen, T., Tanser, F., Gqwede, Z., Mbizana, C., Herbst, K., Newell, M.-L., 2008. High HIV incidence in a community with high HIV prevalence in rural South Africa: findings from a prospective population-based study. AIDS 22 (1), 139–144.

Bartoli, F., 2014. Attitudes Towards Information and Beliefs: How Do Biased Perceptions Influence Individuals' Behaviour? Mimeo. University of Siena.

Beauclair, R., Delva, W., 2013. Is younger really safer? A qualitative study of perceived risks and benefits of age-disparate relationships among women in Cape Town, South Africa. PLoS ONE 8 (11), e81748, http://dx.doi.org/10.1371/journal.pone.0081748.

Becker, G., 1974. A theory of marriage. In: Schultz, T.W. (Ed.), Economics of the Family. University of Chicago Press, Chicago, pp. 299–344.

Bertsch, S., Pesta, B.J., Wiscott, R., McDaniel, M.A., 2007. The generation effect: a meta-analytic review. Mem. Cognit. 35 (2), 201–210.

[25] Note that while this paper and Dupas (2011) share underlying diagnosis for the prevalence of age-disparate sex, the outcomes measured are different: Dupas (2011) measures effects on behavioral and health outcomes but does not measure effects on beliefs, which this paper focuses on. As such, no direct comparison of effect sizes is possible.

[26] All schools with Grade 8–12 students are equipped with computer laboratories.

[27] See Supplementary Material for details on these calculations.

Buseh, A., Glass, L., McElmurry, B., 2002. Cultural and gender issues related to HIV/aids prevention in rural Swaziland: a focus group analysis. Health Care Women Int. 23 (2), 173–184.

Cluver, L., Boyes, M., Orkin, M., Pantelic, M., Molwena, T., Sherr, L., 2013. Child-focused state cash transfers and adolescent risk of HIV infection in South Africa: a propensity-score-matched case–control study. Lancet Global Health 1 (6), e362–e370.

Dupas, P., 2011. Do teenagers respond to HIV risk information? Evidence from a field experiment in Kenya. Am. Econ. J.: Appl. Econ. 3 (1), 1–34.

El-Gamal, M., Grether, D., 1995. Are people Bayesian? Uncovering behavioural strategies. J. Am. Stat. Assoc. 90 (December), 432.

Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10 (2), 171–178.

Garnett, G.P., Anderson, R.M., 1993. Factors controlling the spread of HIV in heterosexual communities in developing countries: patterns of mixing between different age and sexual activity classes. Philos. Trans. R. Soc. Lond. Ser. B: Biol. Sci. 342 (1300), 137–159.

Glennerster, R., Takavarasha, K., 2010. Empowering Young Women: What Do We Know? Technical Report. Originally Prepared for the Nike Foundation by The Abdul Latif Jameel Poverty Action Lab at MIT.

Glynn, J., Carael, M., Auvert, B., Kahindo, M., Chege, J., Musonda, R., Kaona, F., Buve, A., 2001. Why do young women have a much higher prevalence of HIV then young men? A study in Kisumu, Kenya and Ndola, Zambia. AIDS 15 (Suppl. 4), S51–S60.

Gregson, S., Nyamukapa, C., Garnett, G., Mason, P., Zhuwau, T., Caraël, M., Chandiwana, S., Anderson, R., 2002. Sexual mixing patterns and sex-differentials in teenage exposure to HIV infection in rural Zimbabwe. Lancet 359 (9321), 1896–1903.

Griffin, D., Tversky, A., 1992. The Weighing of Evidence and the Determinants of Confidence. Cognit. Psychol. 24, 411–435.

Harling, G., Newell, M.L., Tanser, F., Kawachi, I., Subramanian, S.V., Barnighausen, T., 2014. Do age-disparate relationships drive HIV incidence in young women? Evidence from a population cohort in rural KwaZulu-Natal, South Africa. J. Acquir. Immun. Defic. Syndr. 66 (4), 443–451.

Kahneman, D., 2011. Thinking, Fast and Slow. Macmillan, New York.

Kahneman, D., Knetsch, J.L., Thaler, R.H., 1991. Anomalies: the endowment effect, loss aversion, and status quo bias. J. Econ. Perspect., 193–206.

Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. Econometrica 47 (2), 263–291.

Katz, I., Low-Beer, D., 2008. Why has HIV stabilized in South Africa, yet not declined further? Age and sexual behavior patterns among youth. Sex. Transm. Dis. 35 (10), 837–842.

Kaufman, M.R., Mooney, A., Kamala, B., Modarres, N., Karam, R., Ng'wanansabi, D., 2013. Effects of the Fataki campaign: addressing cross-generational sex in Tanzania by mobilizing communities to intervene. AIDS Behav. 17 (6), 2053–2062.

Kelly, R.J., Sewankambo, N.K., Serwadda, D., Wabwire-Mangen, F., Lutalo, T., Wawer, M.J., 2003. Age differences in sexual partners and risk of HIV-1 infection in rural Uganda. J. Acquir. Immun. Defic. Syndr. 32 (4), 446–451.

Kohler, H.-P., Thornton, R.L., 2012. Conditional cash transfers and HIV/AIDS prevention: unconditionally promising? World Bank Econ. Rev. 26 (2), 165–190.

Langeni, T., 2007. Contextual factors associated with treatment-seeking and higher-risk sexual behaviour in Botswana among men with symptoms of sexually transmitted infections. Afr. J. AIDS Res. 6 (3), 261–269.

Leclerc-Madlala, S., 2003. Transactional sex and the pursuit of modernity. Soc. Dyn. 29 (2), 213–233.

Leclerc-Madlala, S., 2008. Age-disparate and intergenerational sex in southern Africa: the dynamics of hypervulnerability. AIDS 22 (Suppl. 4), S17–S25.

List, J.A., 2003. Does market experience eliminate market anomalies? Q. J. Econ. 118 (1), 41–71.

Longfield, K., et al., 2004. Relationships between older men and younger women: implications for STIs/HIV in Kenya. Stud. Fam. Plan. 35 (2), 125–134.

Luke, N., 2003. Age and economic asymmetries in the sexual relationships of adolescent girls in sub-Saharan Africa. Stud. Fam. Plan. 34 (2), 67–86.

Luke, N., 2005. Confronting the sugar daddy stereotype: age and economic asymmetries and risky sexual behavior in urban Kenya. Int. Fam. Plan. Perspect. 31 (1), 6–14.

Maughan-Brown, B., Kenyon, C., Lurie, M.N., 2014. Partner age differences and concurrency in South Africa: implications for HIV-infection risk among young women. AIDS Behav., http://dx.doi.org/10.1007/s10461-014-0828-6.

Nattrass, N., Maughan-Brown, B., Seekings, J., Whiteside, A., 2012. Poverty, sexual behaviour, gender and HIV infection among young black men and women in Cape Town, South Africa. Afr. J. AIDS Res. 11 (4), 307–317.

Pettifor, A., Kleinschmidt, I., Steffenson, A., Macphail, C., Hlongwa-Madikizela, L., Vermaak, K., Padian, N., 2005. Young people's sexual health in South Africa: HIV prevalence and sexual behaviors from a nationally representative household survey. AIDS 19 (14), 1525–1534.

Poulin, M., 2007. Sex, money, and premarital partnerships in southern Malawi. Soc. Sci. Med. (1982) 65 (11), 2383–2393.

Royce, R.A., Seña, A., Cates, W., Cohen, M.S., 1997. Sexual transmission of HIV. N. Engl. J. Med. 336 (15), 1072–1078, http://dx.doi.org/10.1056/NEJM199704103361507.

Sharp, C., Dellis, A., 2010. Risky Teens: Intervention Science and Its Application to the South African Context. Report Commissioned by the South African Responsible Gambling Foundation.

Shisana, O., Rehle, T., Simbayi, L., Mbelle, N., 2005. South African National HIV Prevalence, HIV Incidence, Behaviour and Communication Survey, 2005. Human Sciences Research Council Press.

Shisana, O., Rehle, T., Simbayi, L., Zuma, K., Jooste, S., Pillay-Van-Wik, V., Mbelle, N., Van Zyl, J., Zungu, N.P., Pezi, Z., Team, S.I.I., 2009. South African National Prevalence, Incidence, Behaviour and Communication Survey, 2008: A Turning Tide Among Teenagers? Human Sciences Research Council Press.

South African Department of Health, 2012. The 2011 National Antenatal Sentinel HIV & Syphilis Prevalence Survey in South Africa. National Department of Health.

Sunstein, C.R., 2005. Precautions against what? The availability heuristic and cross-cultural risk perception. Ala. Law Rev., 57–75.

Tversky, A., Kahneman, D., 1973. Availability: A Heuristic for Judging Frequency and Probability. Cognit. Psychol. 5, 207–232.

UNAIDS, 2012a. UNAIDS Report on the Global AIDS Epidemic, Technical Report. UNAIDS.

UNAIDS, 2012b. UNAIDS World AIDS Day Report 2012. UNAIDS, Geneva.

UNAIDS, 2013a. 2013 Global Fact Sheet. UNAIDS, Geneva, Available from: http://www.unaids.org/en/media/unaids/contentassets/documents/epidemiology/2013/gr2013/20130923_FactSheet_Global_en.pdf

UNAIDS, 2013b. Global Report: UNAIDS Report on the Global AIDS Epidemic 2013. UNAIDS, Geneva.

Varghese, B., Maher, J.E., Peterman, T.A., Branson, B.M., Steketee, R.W., 2002. Reducing the risk of sexual HIV transmission: quantifying the per-act risk for HIV on the basis of choice of partner, sex act, and condom use. Sex. Transm. Dis. 29 (1), 38–43.

de Walque, D., Dow, W.H., Nathan, R., Abdul, R., Abilahi, F., Gong, E., Isdahl, Z., Jamison, J., Jullu, B., Krishnan, S., et al., 2012. Incentivising safe sex: a randomised trial of conditional cash transfers for HIV and sexually transmitted infection prevention in rural Tanzania. Br. Med. J. Open 2 (1).