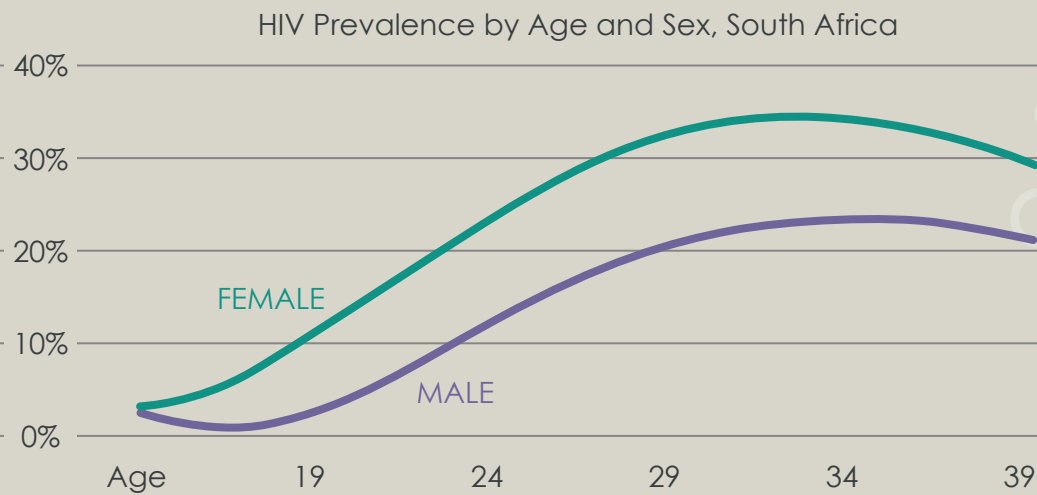


Risking It All for Love?



Resetting Beliefs About HIV Risk Among Low-Income South African Teens

Authors:

Saugato Datta, *ideas42*

Justine Burns, *University of Cape Town*

Matthew Darling, *ideas42*

Katherine Eyal, *University of Cape Town*

Brendan Maughan-Brown, *University of Cape Town*

2014

Working Paper

Abstract¹

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”, which provides repeated doses of information about the link between age and HIV-risk leads to better understanding of this relationship among low-income South African adolescents than a traditional “brochure approach” to dispensing information about HIV risk does. Our results are striking: our randomly-assigned treatment group 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. Subjects in the treatment group answer, on average, twice as many questions about HIV risk and age correctly than those in the control group. We also find that subjects’ beliefs about HIV risk among women is much less accurate than their beliefs about HIV risk among men. Finally, a follow-up survey shows substantially higher retention among Treatment subjects than among Control subjects, though we cannot be sure that this is due to differentially biased selection into the pool of follow-up subjects.

¹ This research was made possible by the generous support of the Department of the Premier and the Department of Health, Western Cape Government, South Africa. Zhanrui Kuang provided excellent support on data analysis. We also thank our experimental cohort and focus group participants.

Table of Contents

I. Introduction.....	4
II. Evidence on Reducing Age-Disparate Partnerships.....	7
III. Changing Beliefs About HIV Risk and Age	8
IV. Experimental Design.....	9
V. Subject Recruitment and Experimental Protocol.....	10
VI. Characteristics of Experimental Sample and Treatment and Control Groups	12
VII. Results.....	13
VII. Discussion.....	19
References.....	20

I. 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 million² 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 2012*).

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, and 50% of those living with HIV in all low- and middle-income countries (*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. 2009; Nattrass et al. 2012; Pettifor et al. 2005; Shisana et al. 2005*).³

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 & Low-Beer 2008; Kelly et al. 2003; Gregson et al. 2002; Garnett & Anderson 1993*). Such relationships are widely prevalent across sub-Saharan Africa, as multiple qualitative and quantitative studies have shown (*Luke 2003*). National survey data on age-mixing among 15-to-24 year old South African women, for example, found that 32.6% of partnerships had an age-gap of at least five years (*Pettifor et al. 2005*). 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,⁴ 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

² Estimates for global prevalence range between 23.5million and 26.6 million.

³ For example, *Shisana et al. (2009)* find that HIV prevalence among South African women under 30 is more than twice that among men under 30. Similarly, 20-30 year old Black African women in the Cape Area Panel Study were three times as likely to have HIV than 20-30 year old Black African men (*Nattrass et al. 2012*). Other studies have found HIV prevalence among 15-24 year old women that are roughly four times those for 15-24 year old men (*Pettifor et al. 2005; Shisana et al. 2005*).

⁴ This survey studied HIV prevalence, HIV incidence, behavior and communication.

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*).⁵

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 were 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*). 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,⁶ being disproportionately driven by recent, salient discussion of deaths due to HIV (*see Sunstein 2005*).

⁵ 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*). 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. 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. 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*). In Botswana, a large study among men found that for each one-year increase in the age-gaps between partners there was a 28% increase in the odds of having unprotected sex (*Langeni 2007*).

⁶ 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.

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 experimental task 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 eight 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 more than doubles the mean number of questions a subject answers correctly. 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”.

The remainder of this paper is organized as follows. In Section II, we review the existing evidence on interventions to tackle age-disparate sex and related behaviors among teens. Section III delves into the literature in behavioral economics 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 IV describes and provides the rationale for the intervention we ran. Section V describes how our experimental subjects were recruited and how the experiment was conducted. Section VI describes our data and the outcome variables we use. Section VII describes our key results. Section VIII concludes with a discussion of these results and implications for further research.

II: 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 towards different classes of interventions. The emphasis on economic benefits points towards the use of cash transfers, both conditional and unconditional, as a way of reducing the attractiveness of age-disparate partnerships. The role of social status and cultural capital point towards 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 (2009), 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.

III. 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, Kahneman and Tversky (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 (ibid.)”. 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 run with adolescents from low income communities in Cape Town as a precursor to our experimental work, we found widespread evidence of incorrect beliefs about the relationship between HIV and age,⁷ beliefs that could plausibly lead young people to choose age-disparate partnerships, with possible repercussions for their own risk of acquiring HIV. Developing effective interventions that rapidly and effectively improve young women and older men’s understanding of the HIV-risks associated with age-disparate partnerships thus emerges as a priority, as noted by Leclerc-Madlala (2008).

⁷ These focus groups delved into subjects’ attitudes towards sexual relationships (age-disparate and otherwise), their beliefs about HIV risk, and the drivers of their sexual behavior. While not intended to be rigorously analyzable, the responses of our teenaged subjects suggested that the phenomenon of an incorrect understanding of the relationship between age and HIV risk was at play in the context we were interested in, i.e. among low-income adolescents in the Cape Town area. 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.

IV. Experimental Design

The purpose of our experimental task 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 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.

In contrast, Treatment subjects played eight rounds of a computer-based “HIV risk game”, programmed and conducted with the software z-Tree (*Fischbacher 2007*).

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.⁸

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 task 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.⁹

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.

⁸ 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.

⁹ Similarly, Berberis, Shleifer and Vishny (1998) and Griffin and Tversky (1992) present evidence in favour of these two heuristics in individual decision making.

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.

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 has been observed in numerous laboratory settings (Kahneman, Knetsch and Thaler, 1990) 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 Berstch 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.

V. Subject Recruitment and Experimental Protocol

The experiment was advertised using posters (see appendix item) at and around a public library 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 form 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 minutes on average.

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 experimental

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

task 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) instructions¹¹ were built into the program itself. In addition, the research assistant conducting the session explained them to participants in both Xhosa and English, and these assistants continued to circulate for the duration to answer questions of clarification.¹² 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 8 rounds of the “HIV risk game” described in Section IV.

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.¹³ 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).¹⁴ This concluded the experiment.

After the experiment, both treatment and control subjects were given an opportunity to click through a series of questions asking them about the relative risk of pairs of individuals, but without any feedback as to the (in)correctness of their answers or any hints. Whoever got the maximum number of correct answers¹⁵ would receive an additional R50 (approximately \$5) payment, in addition to the R50 participation fee received by all subjects. While this was not part of our experiment as such, it was done to give subjects an additional incentive to be attentive during the main experiment.

Furthermore, approximately 3 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.

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 Note that the experiment was not timed.

¹³ Subjects also answered questions about the prevalence of and attitudes towards 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.

¹⁴ 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 peak prevalence age is 30-34 for men and 25-29 for women, this suggested that limiting the age comparisons of the hypothetical individuals in the task 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.

¹⁵ In practice, this usually meant the individual or individuals who got all answers correct.

VI. Characteristics of Experimental Sample and Treatment and Control Groups

Our sample consists of 162 individuals recruited as described in Section V. Table 1 summarizes the information collected on the experimental sample's age, ethnicity, and gender.¹⁶ As it shows, 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.¹⁷ 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.¹⁸ 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.

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 characteristics—age, gender, and race/ethnicity—between those in the treatment and control groups.¹⁹ 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.

Table 2: Characteristics of Individuals in Treatment and Control Groups

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

1. Numbers in parentheses are standard errors

2. Stars indicate statistical significance: *= $p < 0.10$, **= $p < 0.05$, ***= $p < 0.01$

3. For age, $n=79$ in treatment and 77 in control as 6 subjects did not enter an exact age.

¹⁶ 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. In addition to this, we collected participants' school grade level at school, which we do not report due to it being highly correlated, as expected, with age, which is recorded accurately in all but 6 cases. In contrast, 15 participants chose "Not in School" for the Grade variable.

¹⁷ We do not accurately know the ages of 6 individuals who chose either "older than 20" or "younger than 15".

¹⁸ 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.

¹⁹ 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 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

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.

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 V, 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.²⁰ 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.

VII: Results

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

Table 3: Percent of Subjects Answering Questions on HIV Risk-Age Relationships Correctly

	Panel A: All Subjects			Panel B: Male Subjects			Panel C: Female Subjects		
	Control	Treatment	Difference (T-C)	Control	Treatment	Difference (T-C)	Control	Treatment	Difference (T-C)
% Correctly Identified Older Man as Riskier	0.63 (0.05)	0.80 (0.04)	0.18** (0.07)	0.52 (0.07)	0.77 (0.06)	0.24** (0.10)	0.77 (0.08)	0.85 (0.06)	0.08 (0.10)
% Correctly Identified Older Woman as Riskier	0.30 (0.05)	0.76 (0.05)	0.45** (0.07)	0.35 (0.07)	0.79 (0.06)	0.43** (0.09)	0.23 (0.08)	0.71 (0.08)	-0.48** (0.11)
n	80	82	162	48	47	95	31	34	65

1. Numbers in parentheses are standard errors
 2. Stars indicate statistical significance: *= $p < 0.10$, **= $p < 0.05$, ***= $p < 0.01$

²⁰ We chose the ages 20 and 30 because we were particularly keen to use ages that would be reasonable approximations to the ages of proximate-age and older partners for our subjects, who were aged 15-20, as well as to ensure that we were asking questions whose answer would be obvious to anyone who had grasped the basic heuristic we were trying to convey, rather than requiring an in-depth knowledge of the precise prevalence rates by cohort. Given that HIV prevalence is available by 5-year cohorts, we settled on the 20 vs 30 comparison as one that fit both these criteria. While we could have asked several further questions which varied the ages slightly, we chose not to after field testing indicated that subjects – particularly control subjects – were taking longer to complete the task than anticipated, likely due to slower reading speeds than expected.

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, 63% of whom answered correctly. This difference is both large in magnitude and statistically significant. The results are qualitatively similar, 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.

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.²¹ The basic specification (Column I and Column IV) has no additional controls; Columns

Table 4: Effect of Treatment on Probability of Correctly Identifying Older Individual as Riskier

Independent Variable	Panel A			Panel B		
	Probability of Correctly Identifying Older Man as Riskier			Probability of Correctly Identifying Older Woman as Riskier		
	I	II	III	IV	V	VI
Treatment	0.18** (0.07)	0.16** (0.07)	0.19** (0.09)	0.46*** (0.07)	0.4*** (0.07)	0.44*** (0.07)
Treatment x Female			-0.07 (0.14)			0.03
Age	N	Y	Y	N	Y	Y
Gender	N	Y	Y	N	Y	Y
Ethnicity	N	Y	Y	N	Y	Y
Session	N	Y	Y	N	Y	Y
n	162	162	162	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
3. Stars indicate significance: * = $p < .10$, ** = $p < .05$, *** = $p < .01$.

²¹ 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.

II and V add fixed effects for demographic variables (age, gender, and race); and Columns III and VI add 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.²²

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 those in the control group. They were also 36 percentage points more likely to get both answers right. Figure 1, which shows the frequency distribution of the number of correct answers for Treatment and Control groups side by side, provides a visual representation of the shift in the distribution of the number of correct answers.


Table 5: Number of Correct Responses to HIV-Age questions

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

Notes:

1. Numbers in parentheses are standard errors
2. Stars indicate statistical significance: *= $p < 0.10$, **= $p < 0.05$, ***= $p < 0.01$

Figure 1: Distribution of Number of Correct Answers, Control v. Treatment



²² 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 task, but rather the nature of the task itself.

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 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 0.47 questions right, this implies that the treatment more than doubles the average number of correct answers.

Table 6: Effect of Treatment on Number of Correct Answers to Age-HIV Risk Questions

Independent Variable	I	II	III
Treatment	0.64*** (0.11)	0.60*** (0.11)	0.57*** (0.15)
Treatment x 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 experimental subject
3. Stars indicate significance: * = $p < .10$, ** = $p < .05$, *** = $p < .01$.

Result 3: Control Subjects’ beliefs about women are less accurate than their beliefs about men

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.

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. A test of proportions shows²³ that the difference between controls’ probability of getting the “male question” right and their probability of getting the “female question” right is statistically

²³ 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.

significant at the 1% level. This result remains robust even when one examines responses for male and female control subjects separately.²⁴ The control group (and particularly females in this group) thus have more inaccurate beliefs about HIV risk among women than they do about HIV risk in men. Alternatively, their beliefs about HIV risk among men are more accurate than their priors about HIV risk in women. These results are not the focus of this study, but we present them here to note that these differences in how individuals believe HIV risk varies with age among men and among women are worthy of further exploration.

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 V, 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.

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

Table 7: Characteristics of Individuals in Follow-up Sample

	Control	Treatment	Difference
Age	17.95 (.23)	17.71 (.26)	.23 (.35)
% Male	.55 (.08)	.56 (.09)	-.01 (.12)
School Grade	11.17 (.20)	11.19 (.17)	-.02 (.27)
% Black	1 (0)	.97 (.03)	.03 (.03)
n	38	32	70

Notes:

1. Numbers in parentheses are standard errors
2. Stars indicate statistical significance: *= $p < 0.10$, **= $p < 0.05$, ***= $p < 0.01$

²⁴ Similarly, 77% of control females correctly identify the older man as more likely to have HIV, but only 23% of them identify the older woman as similarly more likely to have HIV. These proportions also fail a test of equality at the 1% level. The same holds true for males in the control group, although this difference is only marginally significant (at the 10% level).

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 (0.08)	0.78 (0.07)	0.18 (0.11)
n	38	32	70
Just male subjects	0.71 (.1)	0.72 (.11)	0.01 (.14)
n	21	18	39
Just female subjects	0.47 (.12)	0.85 (.10)	0.38** (.16)
n	17	13	30

Notes:

1. Numbers in parentheses are standard errors
2. Stars indicate statistical significance: *= $p < 0.10$, **= $p < 0.05$, ***= $p < 0.01$

persistent effects on understanding of HIV risk. However, note that take-up of our “follow-up survey” is clearly not random. In other words, those who did respond to our surveyors were not randomly selected from among our original participants. In particular, the results in Table 8 might 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 the HIV, 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 excluded from our follow-up survey. 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

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	.55 (.08)	.71 (.07)	-.16 (.11)
% Got woman question right in lab	.26 (.07)	.34 (.08)	-.08 (.10)
# of Questions Correct	.81 (.12)	1.05 (.12)	-.24 (.17)
n	42	38	80

Notes:

1. Numbers in parentheses are standard errors
2. Stars indicate statistical significance: *= $p < 0.10$, **= $p < 0.05$, ***= $p < 0.01$

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	.76 (.06)	.88 (.06)	-.12 (.08)
% Got woman question right in lab	.72 (.06)	.81 (.07)	-.09 (.09)
# of Questions Correct	1.48 (.09)	1.69 (.10)	-.21 (.14)
n	50	32	82

Notes:

1. Numbers in parentheses are standard errors
2. Stars indicate statistical significance: $*=p<0.10$, $**=p<0.05$, $***=p<0.01$

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).

While we cannot tell whether this bias was larger in the case of the treatment group, which is what really matters for our interpretation of the differences documented in Table 8, its existence suggests that we should treat our follow-up results with a degree of caution. In other words, we are not in a position to work out whether the differences we see in Table 8 are truly the result of greater recall among female Treatment group members, or whether the result is driven by differential selection into the follow-up sample from among treatments and controls. Nonetheless, we are encouraged by the extent of the difference between treatment and control females' ability to correctly answer the HIV-risk question they were asked three months after they were exposed to the information in our intervention.

VIII. 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. 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 – albeit possibly driven by selection – 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” the task of information acquisition may be a far more effective way to deliver information about matters—such as HIV risk—which it 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 (2009) 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. Condom use also increased, and teenaged girls in the treatment group reported younger sexual partners than previously. The prototype game we test here conveys much the same information as the classroom education in Dupas (2009) 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.

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 effects, if any, on attitudes towards age-disparate relationships, the propensity to engage in them, and wider effects on risky sexual behavior.

References:

- Anderson, K. G. & Beutel, A. M. (2007), 'HIV/AIDS prevention knowledge among youth in Cape Town, South Africa', *Journal of Social Sciences* 3(3), 143–150.
- Anderson, K. G., Beutel, A. M. & Maughan-Brown, B. (2007), 'HIV risk perceptions and first sexual intercourse among youth in Cape Town, South Africa', *International Family Planning Perspectives* pp. 98–105.
- Airhihenbuwa, C., Okoror, T., Shefer, T., Brown, D., Iwelunmor, J., Smith, E., Adam, M., Simbayi, L., Zungu, N., Dlakulu, R. et al. (2009), 'Stigma, culture, and HIV and AIDS in the Western Cape, South Africa: An application of the pen-3 cultural model for community-based research', *Journal of Black Psychology* 35(4), 407–432.
- 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 Series*.
- 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. *African Journal of Reproductive Health*, 11(3), pp.197–220.
- Bärnighausen, Till, Frank Tanser, Zanomsa Gqwede, Clifford Mbizana, Kobus Herbst, and Marie-Louise Newell. "High HIV incidence in a community with high HIV prevalence in rural South Africa: findings from a prospective population-based study." *Aids* 22, no. 1 (2008): 139-144.
- Bartoli, F (2014) Attitudes towards information and beliefs: how do biased perceptions influence individuals' behaviour?"
- Beauclair, R., Temmerman, M., Welte, A., and Delva, W. 2012. Age-disparate relationships and implications for STI transmission among young adults in Cape Town, South Africa. *The European Journal of Contraception and Reproductive Health Care*, 17(1), pp.30–39.
- Becker G., 1974, "A theory of marriage" in T.W. Schultz (ed.) *Economics of the Family*, Chicago-London, University of Chicago Press, pp. 299-344.
- Bertsch, S., Pesta, B. J., Wiscott, R., & McDaniel, M. A. (2007). The generation effect: A meta-analytic review. *Memory & Cognition*, 35(2), 201-210.
- 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 for Women International*, 23(2), pp.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. *The Lancet Global Health*, 1(6), e362-e370.
- Dupas, P. (2005), 'Relative risks and the market for sex: Teenagers, sugar daddies and HIV in Kenya'.

Dupas, P. (2009), Do teenagers respond to HIV risk information? Evidence from a field experiment in Kenya, Technical report, National Bureau of Economic Research.

El-Gamal, M. and Grether, D. (1995) "Are People Bayesian? Uncovering Behavioural Strategies", *Journal of the American Statistical Association*, Dec, Vol. 90, No. 432

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. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 342(1300), pp.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. and Buve, A. 2001. Why do young women have a much higher prevalence of HIV than young men? A study in Kisumu, Kenya and Ndola, Zambia. *AIDS*, 15(supple 4), pp.S51–S60.

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

Hope, R. (2007), 'Addressing cross-generational sex: a desk review of research and programs.

Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. *The journal of economic perspectives*, 193-206.

Kahneman, D. & Tversky, A., 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), pp. 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), pp.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 and Behavior* 17(6), 2053–2062.

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

Kohler, H.-P. & Thornton, R. L. (2012), 'Conditional cash transfers and HIV/AIDS prevention: Unconditionally promising?', *The World Bank Economic Review* 26(2), 165–190.

Kong, X., Kigozi, G., Nalugoda, F., Musoke, R., Kagaayi, J., Latkin, C., Ssekubugu, R., Lutalo, T., Nantume, B., Boaz, I. et al. (2012), 'Assessment of changes in risk behaviors during 3 years of post-trial follow-up of male circumcision trial participants uncircumcised at trial closure in Rakai, Uganda', *American Journal of Epidemiology* 176(10), 875–885.

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), pp.261–269.

Leclerc-Madlala, S., 2003. Transactional sex and the pursuit of modernity. *Social Dynamics*, 29(2), pp.213–233.

Leclerc-Madlala, S., 2008. Age-disparate and intergenerational sex in southern Africa: the dynamics of hypervulnerability. *AIDS*, 22 Suppl 4, pp.S17–25.

List, J. A. (2003). Does market experience eliminate market anomalies?. *The Quarterly Journal of Economics*, 118(1), 41-71.

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

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

Luke, N., 2005. Confronting the “sugar daddy” stereotype: age and economic asymmetries and risky sexual behavior in urban Kenya. *International Family Planning Perspectives*, 31(1), pp.6–14.

Mall, S., Middelkoop, K., Mark, D., Wood, R. & Bekker, L.-G. (2013), ‘Changing patterns in HIV/AIDS stigma and uptake of voluntary counselling and testing services: The results of two consecutive community surveys conducted in the Western Cape, South Africa’, *AIDS care* 25(2), 194–201.

Mitchell, K. J., Bull, S., Kiwanuka, J. & Ybarra, M. L. (2011), ‘Cell phone usage among adolescents in Uganda: acceptability for relaying health information’, *Health Education Research* 26(5), 770–781.

Nattrass, N., Maughan-Brown, B., Seekings, J., and 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), pp.307–317.

Nglazi, M. D., van Schaik, N., Kranzer, K., Lawn, S. D., Wood, R. & Bekker, L.-G. (2012), ‘An incentivized HIV counseling and testing program targeting hard-to-reach unemployed men in Cape Town, South Africa’, *J AIDS Journal of Acquired Immune Deficiency Syndromes* 59(3), e28–e34.

Ott, M. Q., Barnighausen, T., Tanser, F., Lurie, M. N. & Newell, M.-L. (2011), ‘Age-gaps in sexual partnerships: seeing beyond “sugar daddies”’, *AIDS (London, England)* 25(6), 861-863.

Pettifor, A., Kleinschmidt, I., Steffenson, A., Macphail, C., Hlongwa-Madikizela, L., Vermaak, K., and 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), pp.1525–1534.

Poulin, M., 2007. Sex, money, and premarital partnerships in southern Malawi. *Social science & medicine* (1982), 65(11), pp.2383–2393.

Royce RA, Seña A, Cates W, Cohen MS. Sexual transmission of HIV. *New England Journal of Medicine*. 1997;336(15):1072–1078. doi:10.1056/NEJM199704103361507.

Sharp, C. & Dellis, A. (2010), 'Risky teens: Intervention science and its application to the South African context', South Africa.

Shisana, O., Rehle, T., Simbayi, L., and Mbelle, N. 2005. *South African National HIV prevalence, HIV incidence, behaviour and communication survey, 2005*, Pretoria: 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. and Team, S.I.I. 2009. *South African National Prevalence, Incidence, Behaviour and Communication Survey, 2008: A Turning Tide Among Teenagers?* Cape Town: 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. L. Rev.*, 57, 75.

Townsend, L., Jewkes, R., Mathews, C., Johnston, L. G., Flisher, A. J. Zembe, Y. & Chopra, M. (2011), 'HIV risk behaviours and their relationship to intimate partner violence (IPV) among men who have multiple female sexual partners in Cape Town, South Africa', *AIDS and Behavior* 15(1), 132–141.

UNAIDS (2012), *UNAIDS report on the global AIDS epidemic*, Technical report, UNAIDS.

UNAIDS, 2012. *UNAIDS World AIDS Day Report 2012*, Geneva: UNAIDS.

UNAIDS, 2013a. *2013 Global Fact Sheet*, Geneva: UNAIDS. Available at: 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*, Geneva: UNAIDS.

Urs Fischbacher (2007): z-Tree: Zurich Toolbox for Ready-made Economic Experiments, *Experimental Economics* 10(2), 171-178.

Varghese B, Maher JE, Peterman TA, Branson BM, Steketee RW. 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* 2002;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', *BMJ open* 2(1).