Published in final edited form as:

J R Stat Soc Ser A Stat Soc. 2009 April ; 172(2): 443–465. doi:10.1111/j.1467-985X.2009.00585.x.

Analyzing Direct Effects in Randomized Trials with Secondary Interventions: An Application to HIV Prevention Trials

Michael Rosenblum[†],
University of California, Berkeley, USA

Nicholas P. Jewell,
University of California, Berkeley, USA

Mark van der Laan,
University of California, Berkeley, USA

Steve Shiboski,
University of California, San Francisco, USA

Ariane van der Straten, and
University of California, San Francisco, USA

Nancy Padian
University of California, San Francisco, USA

Summary

The Methods for Improving Reproductive Health in Africa (MIRA) trial is a recently completed randomized trial that investigated the effect of diaphragm and lubricant gel use in reducing HIV infection among susceptible women. 5,045 women were randomly assigned to either the active treatment arm or not. Additionally, all subjects in both arms received intensive condom counselling and provision, the “gold standard” HIV prevention barrier method. There was much lower reported condom use in the intervention arm than in the control arm, making it difficult to answer important public health questions based solely on the intention-to-treat analysis. We adapt an analysis technique from causal inference to estimate the “direct effects” of assignment to the diaphragm arm, adjusting for condom use in an appropriate sense. Issues raised in the MIRA trial apply to other trials of HIV prevention methods, some of which are currently being conducted or designed.

Keywords

Causal inference; Intention-to-treat; Randomized trials; Time-dependent confounding

1. Introduction

Randomized controlled trials play a critical role in the assessment of public health interventions to prevent and treat disease. Standard intention-to-treat analyses, which are straightforward comparisons of aggregate outcomes in the treatment and control groups, are widely assumed to be the most appropriate method for analyzing and reporting trial results. However, in intervention trials, ethical considerations generally require provision of currently accepted disease prevention methods to all participants as a secondary intervention, irrespective of study

[†]Address for Correspondence: Michael Rosenblum, CAPS, University of California, San Francisco, 50 Beale St., Suite 1300, San Francisco, CA, 94105 USA. mrosenblum@csail.mit.edu.

arm assignment. For example, in randomized controlled trials of new HIV prevention methods, provision of condoms with counselling is required for all study participants to meet human subjects protection requirements. In such cases, the standard intention-to-treat analysis may not adequately answer important questions of public health interest. In particular, in these cases it may not address the effectiveness of the new prevention method in the absence of the secondary intervention. Also, it may give severely biased estimates of study product efficacy, even in blinded trials with relatively high rates of compliance to the primary intervention (Trussell and Dominik, 2005). The situation is even worse in unblinded trials of new HIV prevention methods; for example, in the MIRA trial of latex diaphragms and gel, the frequency of condom use was different in the treatment and control arms, making it quite difficult to draw conclusions about diaphragm and gel efficacy based only on the intention-to-treat analysis (Padian et al., 2007).

In this paper, we highlight limitations in relying solely on standard analyses of randomized trials with secondary interventions such as condom counselling, and we propose a supplemental analysis technique. Our intent is to stimulate debate on additional design and analysis approaches. Improved methods tailored to HIV prevention trials are urgently needed, since many of these trials are currently being conducted or planned. We focus on supplemental analysis approaches throughout this paper, but briefly consider design issues in the discussion section.

We describe a causal inference technique, called a direct effects analysis, that attempts to isolate the effects of the secondary intervention, thereby providing more information about primary treatment efficacy. This technique can be applied in situations where a secondary intervention affects the primary outcome only through a measured covariate or set of covariates. For example, in HIV prevention trials, it is thought that condom provision and counselling affect HIV risk primarily through their impact on condom use, which is often measured in these trials. Unlike standard regression techniques to adjust for time-dependent confounders, the direct effects analysis adequately handles time-dependent confounders measured in a trial. We illustrate the benefits and limitations of our direct effects analysis by applying it to the trial entitled “Methods for Improving Reproductive Health in Africa” (MIRA) (Padian et al., 2007), a randomized controlled trial investigating the effect of providing diaphragm and non-contraceptive, lubricant gel (Replens), on male-to-female HIV transmission.

We focus primarily on unblinded trials (where participants know which study arm they are in), since here the problems in interpreting the intention-to-treat estimator can be worse than in blinded trials. Unblinding may be a necessary part of a trial design, for example in trials of barrier methods or male circumcision where no placebo is used; some unintended unblinding may also occur in trials with placebo (Friedman et al., 1998). In the MIRA trial, which is unblinded, the fact that the HIV infection rates in the two study arms were nearly identical does not necessarily mean that latex diaphragms offer no protection against HIV; an alternative explanation, whose plausibility we explore in this paper, is that being assigned to the diaphragm arm resulted in significantly less condom use, which may have cancelled out a potentially protective effect from diaphragms. A direct effects analysis aims to isolate the effect of a suspected alternative causal pathway connecting treatment assignment to outcome, in order to help assess whether such alternative explanations are plausible. Our approach is based on a method of direct effects analysis from Robins and Greenland (1992) and Pearl (2000b). This method is also discussed in Petersen et al. (2006), under a different set of assumptions. Though we concentrate on unblinded trials with secondary interventions, we also discuss analogous issues in blinded trials with secondary interventions in the discussion section of the paper.

Our paper is organized as follows: We describe the MIRA trial in Section 2. In Section 3 we discuss the limitations of intention-to-treat analyses in the presence of secondary interventions

or unblinding. In Section 4 we introduce direct effects, and in Sections 5 and 6 we present estimators of direct effects, and the assumptions needed for these estimators to be valid. In Section 7, we provide summary data from the MIRA trial, and then examine the results of standard intention-to-treat analyses as well as direct effects analyses. Lastly, in Section 8 we interpret the results of the direct effects analysis, discuss its advantages and limitations, and then consider alternative analysis and design approaches. We note that there are many ethical issues associated with conducting HIV prevention trials such as the MIRA trial; we do not address these in this paper.

2. The MIRA Trial

We briefly describe the MIRA trial (Padian et al., 2007), used throughout the paper to illustrate the application and interpretation of a direct effects analysis in HIV prevention trials. The goal of the MIRA trial was to determine if latex diaphragms and gel provide protection against HIV infection. The motivation for this research was to find a safe, effective method that could be used by women whose partners would not wear condoms.

The MIRA trial is a randomized controlled trial with two arms. In the “diaphragm arm,” each subject was instructed to use a fitted Ortho All Flex diaphragm, Replens non-contraceptive, lubricant gel, and condoms at all sex acts, and were supplied with these products. In the “control arm,” each subject was instructed to use condoms at all sex acts, and was supplied with condoms. Additionally, in both arms, subjects were given intensive condom counselling, as well as diagnosis and treatment of sexually transmitted infections. The eligibility criteria for enrolment in the study included being a woman 18-49 years old, HIV-negative at baseline, sexually active, and non-pregnant. 5,045 women were enrolled at three sites in South Africa and Zimbabwe from September 2003 until September 2005. Participants were assessed at quarterly clinic visits with a goal of following each participant for two years. The median follow-up period was 21 months. For a detailed description of subject recruitment and retention, eligibility, survey instruments and data collection, see Padian et al. (2007).

Before we looked at any HIV outcome data, we chose as measure of a subject's condom use over the period since her last visit to be the participant's response to the question: “Did you use a condom at your last sexual intercourse?” This choice was based in part on a growing literature regarding accurate measures of condom use based on self-reporting questionnaires (Anderson et al., 1998). Reported condom use at last sex, recorded at each follow-up visit, averaged 53.5% across all visits in the diaphragm arm and 85.1% in the control arm. (Note that other data show an increase in overall condom uptake after recruitment in both arms—see Padian et al. (2007)). Diaphragm use, based on self-reports of use at last sex at each clinic visit, averaged 73% in the diaphragm arm and only 0.15% in the control arm.

To motivate the analysis in the subsequent sections, consider the following very crude adjustment for condom use, where we estimate the number of infections that would have been prevented had condom use been increased in the diaphragm arm to the level observed in the control arm. That is, we estimate the number of infections that would have been prevented, had an additional $(85\% - 53.5\%) = 31.6\%$ of the 2472 subjects in the diaphragm arm used condoms consistently; 31.6% of 2472 subjects is about 781 subjects. First, we calculate how many infections we would expect among 781 randomly chosen subjects in the diaphragm arm who infrequently used condoms. Based on the observed infection rate in the diaphragm arm among subject-visits in which women reported no condom use at last sex (4.2% new infections per woman-year), and the mean length of time subjects were followed-up (1.6 years), we would expect $(781 \times 1.6) \text{ woman-years} \times 4.2\% \text{ new infections per woman-year} \approx 51$ of these 781 subjects to become infected with HIV by the end of the trial. Since consistent condom use is believed to reduce HIV risk by 80% (Weller and Davis-Beatty, 2002), our crude analysis

indicates that we would expect about $80\% \times 51 \approx 41$ prevented infections, had condom use been equalized in the two arms. The number of such prevented infections required to achieve a statistically significant difference between the two study arms, at significance level $\alpha = 0.05$ (two-tailed), is 39. Thus, crudely adjusting for the number of infections attributed to the observed difference in reported condom use between arms leads to a statistically significant difference in infections between the arms, though just barely. This suggests that a more sophisticated analysis at least has the potential to show a protective effect of diaphragms. The direct effects analysis aims to do just this, by adequately dealing with confounders and the time-dependent nature of the data.

A key public health question is how much, if any, protection diaphragms provide against male-to-female HIV infection in a community—as compared to a research—setting. Another important question is how much protection consistent diaphragm and gel use (but no condom use) provide, compared to unprotected sex. As further discussed in Section 3, neither of these public health questions is answered adequately by a standard intention-to-treat comparison of HIV infection rates between study arms.

3. Intention-To-Treat Analysis

The simplest intention-to-treat analysis of a randomized controlled trial compares the average outcome for the treatment group to that for the control group. For example, in the MIRA trial, this would be a comparison of HIV incidence between the diaphragm arm and control arm. The advantages of an intention-to-treat analysis in many types of randomized controlled trials include:

- a. An intention-to-treat analysis is completely protected from confounding, since it does not involve any measurements made after randomization (e.g. adherence to treatment) except for the outcomes of interest.
- b. An intention-to-treat analysis unequivocally answers a specific causal question: specifically, it estimates the effect of the intervention given to the treatment arm compared to the control arm, in the study setting. We refer to the relative risk of infection between treatment and control arms in a particular setting (e.g. in a community setting, in a study setting) as the treatment “effectiveness.”
- c. An intention-to-treat analysis gives a conservative approximation of treatment efficacy. Treatment efficacy, the effect one would observe if treatment were given in an ideal setting where full adherence were guaranteed, is likely to be greater than the effect measured under non-ideal conditions of the trial where there is non-compliance.

Limitations of the intention-to-treat analysis have been pointed out by others; see for example (Sheiner et al., 1995; Frangakis and Rubin, 1999; Hirano et al., 2000; Feinstein, 1985; Friedman et al., 1998; Trussell and Dominik, 2005). We highlight situations in which the intention-to-treat estimator may not answer the most important public health questions related to the primary intervention, when there is a concurrent, secondary intervention or unblinding. First, the simple intention-to-treat comparison may differ substantially from the same comparison in other populations whose members are both subject to different information on the possible efficacy of diaphragms, and who are not provided with free condoms coupled with intensive counselling. Also, while it is well known that in blinded trials, standard intention-to-treat analyses may underestimate product efficacy (Feinstein, 1985; Sheiner et al., 1995; Friedman et al., 1998), these analyses may be even less informative in the presence of a secondary intervention such as condom provision and counselling (Trussell and Dominik, 2005). This can occur, for example, when the secondary intervention differentially affects adherers and non-adherers to the primary study product (Trussell and Dominik, 2005). In short, in the

presence of secondary interventions, the question definitively addressed by the standard intention-to-treat analysis may not be of major public health interest.

In unblinded trials, the intention-to-treat analysis may be severely impacted when the primary intervention results in unintended behaviour changes. For example, in the MIRA trial, those randomized to the diaphragm arm reported much less condom use, as compared to the control arm. Without considering the impact of this differential condom use, an intention-to-treat result of no difference in HIV infections between study arms would be consistent with both (i) the diaphragm being not efficacious, but also (ii) the diaphragm providing some protection that may have been cancelled out by additional infections due to decreased condom use in the diaphragm arm. Unintended effects of study arm assignment are illustrated in the causal diagram of Figure 1 which shows schematically how the overall causal effect of study arm assignment can be viewed as having two distinct causal pathways: (i) a direct effect of treatment on HIV infection (the top arrow), and (ii) an indirect effect through condom use that, in turn, affects the risk of infection (the indirect effect of treatment assignment). Causal diagrams represent sets of assumed causal relationships: for an introduction to causal diagrams and their interpretation see Pearl (2000a) or Chapter 8.2 of Jewell (2004).

4. Direct Effects Analysis: Overview, Definitions and Causal Interpretation

Given the large difference in reported condom use between the two study arms in the MIRA trial (53.5% in the diaphragm arm vs. 85.1% in the control arm), we would like to better understand the role of condom use in mediating the effect of treatment assignment on HIV outcome. For the MIRA trial, we decided to focus on condom use as a mediator for two reasons. First, the condom provision and intensive counselling intervention was believed primarily to affect condom use, and not to affect other HIV risk factors, and it is of major public health interest what the effect of diaphragms would be in the absence of condom counselling, where condom use is generally quite low. Second, except for diaphragm use, condom use was the only time-dependent variable measured in the MIRA trial that was meaningfully different in the two study arms. Variables such as number of other partners and frequency of sex were measured but each showed less than a 5% difference between the two arms.

4.1. Direct Effects Defined through Hypothetical Randomized Trials

Ideally, we would like to know what the infection rates would be in a hypothetical randomized trial in which the diaphragm intervention were given to one arm, and in which all participants were constrained to use condoms infrequently (where we define “infrequently” below). Such a trial is highly unethical so cannot be conducted. In our direct effects analysis, we attempt to approximate the results of such a trial, using data from the MIRA study. If we are successful, this would shed light on the efficacy of diaphragms. For example, if we conclude that in this hypothetical trial there would be no benefit of the diaphragm intervention, this would be strong evidence that the diaphragm was not efficacious in preventing HIV infections (unless actual diaphragm use is much lower than what subjects reported). Alternatively, if we would conclude that in such a trial there would be a strong benefit of the diaphragm intervention, then this would be evidence in favour of diaphragm efficacy.

We describe several hypothetical scenarios below, in which we talk about how many new HIV infections would have occurred had all study subjects been constrained to either infrequently use condoms or consistently use condoms. It has been argued that in order for such hypothetical scenarios to be well defined, it must be possible, at least in principle, to imagine interventions that would lead to such constraints being enforced (Frangakis and Rubin, 2002). We think the constraint on infrequently using condoms could be approximately established by a hypothetical (unethical) intervention. Similarly, in theory, an intervention is possible that would result in condoms consistently being used.

Definition of Direct Effects Relative Risks—Consider what the infection rate would be in the following two hypothetical scenarios: (1) all subjects in the MIRA trial are assigned to the diaphragm arm and are constrained to use condoms consistently; (2) all subjects in the MIRA trial are assigned to the control arm and are constrained to use condoms consistently. We call the ratio of infection rates in the two scenarios the “direct effects relative risk” of treatment assignment with condom use fixed at “consistently.” The ratio of infection rates in the same pair of hypothetical scenarios, except with all subjects constrained to use condoms infrequently, gives the direct effects relative risk with condom use fixed at “infrequently.” This definition of direct effects relative risk corresponds to Type I (“controlled”) direct effects in Petersen et al. (2006). The main challenge in a direct effects analysis is to estimate these quantities using data collected from a randomized trial in which subjects themselves (or their partners) decided when to use condoms, and where condom use is not objectively measured but is self-reported.

4.2. Direct Effects in terms of Potential Outcomes

The above definitions can be written in terms of potential outcomes. The potential outcomes viewpoint has a long history related to work by Fisher and Neyman (as described in Angrist et al. (1996)), and later extended by Rubin (1974) and Robins (1986), though it is not universally accepted (see e.g. Dawid (2000)). In this setting, a potential outcome I_{rc} is defined for each subject in the population under each of the four possible combinations of $r \in \{0, 1\}$, $c \in \{0, 1\}$. Let $r = 0$ indicate the control arm, and $r = 1$ indicate the diaphragm arm; similarly, $c = 0$ indicates infrequently using condoms, and $c = 1$ indicates consistently using condoms. For a given subject, I_{rc} represents whether that subject would have become infected with HIV within 2 years of enrollment in the study, had she been assigned to study arm r and been constrained to use condoms at level c . I_{rc} can either be 0 indicating no HIV infection, or 1 indicating HIV infection. The direct effects analysis involves estimating the probability of HIV infection, $p_{rc} := \Pr(I_{rc} = 1)$ for a randomly chosen subject from the MIRA study, would she have been assigned to arm r and simultaneously constrained to use condoms at fixed level c , for $r \in \{0, 1\}$ and $c \in \{0, 1\}$. These probabilities can then be combined to make appropriate causal comparisons; for example, we can compute the ratio p_{1c}/p_{0c} to get the direct effects relative risk for condom use fixed at level c .

We now describe how the potential outcomes $\{I_{rc}\}$ connect to the observed data in the MIRA trial. For each study subject, the study arm assignment she received in the trial is denoted by R , with $R = 0$ indicating the control arm and $R = 1$ indicating the diaphragm arm. For the time being we assume that use of condoms is simply recorded as $C = 1$ for consistent use and $C = 0$ for infrequent use. We discuss more precisely what we mean by “consistent” and “infrequent”, and how measurement of C was done in the MIRA trial in Section 6. Let I indicate HIV status at the end of the trial, with $I = 0$ indicating not infected, and $I = 1$ indicating infected. We emphasize the distinction between potential (generally unobserved) outcomes I_{rc} and the observed quantities (R, C, I) for each subject. The “consistency assumption” connects the observed data (including R, C, I) to the set of potential outcomes $\{I_{rc}\}$. It states that the observed infection status I is equal to the potential outcome I_{rc} when observed study arm assignment $R = r$ and observed condom use $C = c$. We consider the other, unrealized potential outcomes for each subject as missing data, and apply causal inference methodology of van der Laan and Robins (2003) accordingly.

It is important to note the asymmetry here between the two variables R and C in that the former refers to assignment to a study arm (but not necessarily use of the study product) and the latter refers to actual use of condoms. Thus, direct effects relative risks do not measure, for example, effectiveness of the diaphragm intervention in a setting in which condom use is merely

encouraged; they measure effectiveness of the diaphragm intervention in a setting in which condom use is constrained at a fixed level.

Direct effects analyses provide valuable information for understanding the efficacy of the primary treatment. Though our direct effects relative risks are not targeted at estimating treatment efficacy, and so give biased estimates of efficacy, they do peel away a layer of distortion that results from using the intention-to-treat estimator as an estimate of efficacy. The direct effects analysis removes distortion caused by both (i) the effect of treatment assignment on HIV infection through changing condom use and (ii) the effect of a secondary treatment on HIV infection, insofar as this is captured through changed condom use. The direct effects estimate may give better (though still biased) estimates of efficacy than an intention-to-treat estimator, when the effects in (i) or (ii) are strong, and the remaining effect of treatment assignment on HIV outcome is due primarily to use of the primary study product. The effects of (i) and (ii) are believed to be strong in the case of the MIRA trial, and all data collected indicate that no other measured HIV risk factors (e.g. number of partners, frequency of sex) were affected by treatment assignment. However, it is quite possible that there are HIV risk factors affected by study arm assignment not captured by the study data; this would result in additional bias when direct effects estimates are used as (conservative) estimates of study product efficacy.

5. Estimating Direct Effects in a Single Time Point Setting

We describe how to estimate direct effects in what we call the “single time point setting,” for the MIRA trial. We categorize all participants, based on their self-reported condom use at last sex at each of the first three visits. “Infrequent users” ($C = 0$) are those who reported condom use at last sex in at most one out of the first three visits, and “frequent users” ($C = 1$) are those who reported condom use at last sex in all of the first three visits. We let I represent HIV status at the end of the trial (which consists of a total of eight visits, though some women were not followed-up for the full eight visits). Participants who became infected with HIV in the first three visits were not included in this single time point analysis; even though the number of HIV infections in each arm by the end of the first three visits was nearly identical (84 in the diaphragm arm, 81 in the control arm), excluding these early infections still could potentially lead to bias.

Since arm assignment is randomized, there cannot be confounders of its effect on subsequent HIV status. However, there likely are confounders of the effect of condom use on subsequent HIV status. Such potential confounders measured in the MIRA trial can be categorized as either baseline factors (denoted by W) or time-dependent measurements (denoted by D). Baseline factors measured included (i) personal characteristics (employment status, living with husband, age, ever pregnant, HSV-2 status); (ii) regular partner's characteristics (partner ever tested HIV-positive, partner's age and employment status, whether subject or her partner makes condom use decision); (iii) sexual habits (number of lifetime male partners, frequency of intercourse, other contraceptive use, whether had sex under influence of alcohol in past 3 months, condom use at last sex); and (iv) study site. Time-dependent confounders included diaphragm use at last sex, sexual abstinence since last visit, HSV-2 incidence, and whether sexual partners were known to have multiple partners during the previous three months.

In Sections 5.1 and 5.2, we describe how to estimate the HIV infection probabilities $p_{rc} = \Pr(I_{rc} = 1)$, for all four combinations of $r \in \{0, 1\}$, $c \in \{0, 1\}$; these estimates can then be combined to estimate direct effects relative risks. To present the simplest case of a direct effects analysis, in Section 5.1 we describe this analysis assuming all potential confounders are measured before randomization. In Section 5.2, we extend the estimation techniques to cover confounding variables that are measured post-randomization and may lie on a causal pathway between study

arm assignment and condom use. In Section 6, we deal with time-dependent measurements, including time-dependent confounding. For all of the discussion below, we assume that the data observed on each subject can be viewed as random draw from a hypothetical, infinite population of subjects (called a “superpopulation”).

5.1. Estimation of Direct Effects When All Confounders Can Be Measured Prior to Randomization

We consider the case in which all confounders of the effect of condom use on HIV status are measured before study arm assignment. Such a setting is depicted in Figure 2 of Section 3 above, in which the variable W denotes pre-treatment (baseline) confounders of the effect of condom use (C) on HIV status (I). The goal is to estimate p_{rc} . When all confounders can be measured prior to randomization, as we are assuming in this section, it is possible to use a simple regression model to estimate p_{rc} . However, the purpose of this section is to introduce methods used in later sections that deal with more complicated situations where simple regression methods can be asymptotically biased. We therefore present an inverse probability of treatment weighted (IPTW) estimator (Robins and Finkelstein, 2000) of p_{rc} ; the class of IPTW estimators is asymptotically unbiased, under weaker assumptions than required for simple regression methods to be asymptotically unbiased, in the more complicated situations encountered in later sections.

Before presenting the IPTW estimator, we state assumptions under which it gives consistent, asymptotically normal estimates of the direct effects relative risks. (Consistency implies the estimator converges in probability to the correct value, and asymptotic normality allows for asymptotically valid confidence intervals.) First, we make the following conditional independence assumption, sometimes referred to as the “randomization assumption” (van der Laan and Robins, 2003),

$$C \{I_{rc}\} | W, R. \quad (1)$$

That is, condom use C is independent of the potential outcomes $\{I_{rc}\}$, conditional on baseline confounders W and study arm assignment R . Intuitively, this means that, within strata having the same study arm R and the same baseline measurements W , the decision to use condoms is not an indicator of underlying HIV risk (e.g. due to having high-risk partners, many partners, or having other STIs). This assumption implies that there are no unmeasured confounders of condom use and HIV infection.

The randomization assumption 1 follows from the structural model interpretation of causal diagrams given in (Pearl, 2000a), in which each variable in the causal diagram is assumed to be an unknown function of its parents and an exogenous random variable. The structural model interpretation of the causal diagram in Figure 2 is that for some unknown functions f_1, f_2, f_3, f_4 , and independent (unmeasured) random variables $\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4$, we have that the observed variables satisfy $W = f_1(\varepsilon_1)$, $R = f_2(W, \varepsilon_2)$, $C = f_3(W, R, \varepsilon_3)$, $I = f_4(W, R, C, \varepsilon_4)$, and that the potential outcomes $I_{rc} = f_4(W, r, c, \varepsilon_4)$, for all r, c . This implies assumption (1) above since given W and R , $\{I_{rc}\}$ is a function only of ε_4 , and so is independent of C .

We make the following assumption, called the experimental treatment assignment (ETA) assumption: For all r, c, w , $\Pr(C = c | R = r, W = w) > 0$. That is, no stratum of treatment assignment and baseline confounders precludes the possibility of using condoms at any specific level. If this assumption is false, direct effects will not be identifiable from the data alone, unless further model assumptions are made. We also make the consistency assumption introduced in Section 4.2.

We can connect the probabilities of specific potential outcomes that we ultimately care about to quantities that can be estimated from the observed data using the following chain of equalities:

$$\begin{aligned}
 \Pr(I_{rc}=1) &= \sum_w \Pr(I_{rc}=1|W=w)\Pr(W=w) \\
 &= \sum_w \Pr(I_{rc}=1|W=w, R=r, C=c)\Pr(W=w) \\
 &= \sum_w \Pr(I=1|W=w, R=r, C=c)\Pr(W=w) \\
 &= \sum_w \frac{\Pr(I=1|R=r, C=c, W=w)}{\Pr(R=r|W=w)\Pr(C=c|W=w, R=r)}.
 \end{aligned} \tag{2}$$

The second equality follows from Assumption (1) and the fact that R is randomized so $R \perp \{I_{rc}\}|W$; the third equality follows from the consistency assumption described in Section 4 above; and the last equality follows by basic rules of probability. The experimental treatment assignment (ETA) assumption insures that all the terms in (2) are well-defined. Note that once we get estimates for the left hand side of this equation, for all values of r, c , we can estimate the direct effects relative risks $\Pr(I_{1c} = 1)/\Pr(I_{0c} = 1)$ for each level of fixed condom use c . We also point out that estimating $\Pr(I_{rc} = 1)$ using (2) is equivalent to solving the following estimating equation:

$$E\left(\frac{1}{\Pr(R|W)\Pr(C|W, R)}(I - m(R, C|\beta))\right) = 0, \tag{3}$$

where $m(r, c|\beta)$ denotes a saturated (non-parametric) model for the unknown probability $\Pr(I_{rc} = 1)$.

To estimate (2), we first estimate its numerator using the empirical distribution of the data. The first term in the denominator of (2), $\Pr(R = r|W = w)$, equals 1/2 since R was randomized. The second term in the denominator can be estimated by the empirical distribution if W is low-dimensional, or using a logistic regression model, for example, if W is high-dimensional. Substituting such estimates for the quantities in the right hand side of (2), and summing over w , we get what is referred to as the inverse probability of treatment weighted (IPTW) estimator for $\Pr(I_{rc} = 1)$. We use the bootstrap, as described in Efron and Tibshirani (1993), to generate confidence intervals for direct effects relative risks based on these estimators. Alternatively, the influence curve for the estimator could be used for statistical inference, as described in van der Laan and Robins (2003).

The argument above assumed that all confounders were measured at baseline, prior to randomization. In the next section, we treat the case in which there are confounders of condom use and HIV status that are themselves affected by study arm assignment.

5.2. Estimation of Direct Effects With Confounding by Causal Intermediates

In the MIRA trial, there are potential confounders of condom use and HIV status that are themselves affected by study arm assignment. Such confounders are called “causal intermediates,” and are denoted by D . This situation is illustrated in the causal diagram of Figure 3, where for clarity, pre-randomization potential confounders W are not shown. For an example of a potential confounder that may be a causal intermediate, consider diaphragm use, defined similarly to condom use for the single time point case above—that is, diaphragm use is a dichotomized summary measure, based on self-reports of use at last sex at the initial three

visits. Diaphragm use may affect the decision to use condoms and it also may affect HIV infection (by possibly acting as a physical barrier); thus, diaphragm use may be a confounder that is a causal intermediate. In general, the set of confounders that are causal intermediates will be a vector D including more variables than just diaphragm use (e.g. frequency of sex, number of partners). For concreteness in our explanations, we assume diaphragm use is the only confounder that is a causal intermediate; however in our statistical analyses, we included all the time-dependent variables listed in Section 5. Note that diaphragm use will be “infrequent” ($D = 0$) for virtually all participants in the control arm, except the few who obtained and used diaphragms on their own.

A single regression (e.g. regressing I on R , D , C) cannot be used to estimate the direct effects of R on I , fixing C , in the presence of confounding by causal intermediate D as in Figure 3 (Rosenbaum, 1984; Robins, 1997; Hernán et al., 2000; Petersen et al., 2006). The problem is that if one controls for such a causal intermediate in the regression analysis, one blocks a causal pathway of interest; in the above diagram, by controlling for diaphragm use D , one would block the causal pathway from study arm R through diaphragm use D to HIV status I , which is an important part of what we are trying to estimate. Alternatively, by not controlling for such a causal intermediate, one's estimates may be biased due to confounding from the causal intermediate. The analysis we outline below adequately handles confounding by measured causal intermediates. An additional advantage of using the IPTW analysis below is that, unlike a simple regression analysis, it will still be asymptotically unbiased under the assumptions below even when there are unobserved confounders influencing both the causal intermediate and the outcome. This is depicted in Figure 4 below, where the variable U is such a confounder.

In the presence of confounding by causal intermediates, the randomization assumption analogous to assumption (1) of Section 5.1 is more complicated. It involves not only potential HIV outcomes I_{rc} under different scenarios for study arm assignment r and condom use c , but also potential outcomes for diaphragm use under different possible study arm assignments. We let D_r represent what an individual's diaphragm use would be, would she have been assigned to study arm r . The randomization assumption analogous to (1) now takes the form,

$$C \{I_{rc}, D_r\} \mid W, R, D. \quad (4)$$

Intuitively, given knowledge of a subject's baseline measurements W , study arm assignment R , and diaphragm use D , her observed condom use gives no additional information as to either (1) her risk of HIV infection would she have been assigned to arm r and would she have used condoms at fixed level c or (2) her diaphragm use were she assigned to the other study arm $R = r$. Assumption 4 follows from the structural model interpretation of Figure 3 (Pearl, 2000a), by similar arguments as given just after (1) in Section 5.1.

In the case here of confounding by causal intermediates, the experimental treatment assignment assumption is that for each possible w , r , c , d , we have $\Pr(C = c \mid R = r, D = d, W = w) > 0$. That is, no strata of treatment assignment, diaphragm use, and baseline confounders precludes the possibility of using condoms at any specific level. The consistency assumption in this setting connects the observed data (including W , R , D , C , I) to the set of potential outcomes $\{I_{rc}, D_r\}$. It states that the observed diaphragm use D is equal to the potential outcome D_r when $R = r$; furthermore, it states that infection status I is equal to the potential outcome I_{rc} when observed study arm assignment $R = r$ and observed condom use $C = c$.

Analogous to the derivation in Section 5.1, the assumption (4), the consistency assumption, and the experimental treatment assignment assumption imply:

$$\Pr(I_r=1)=\sum_{w,d}\frac{\Pr(I=1, R=r, C=c, D=d, W=w)}{\Pr(R=r|W=w)\Pr(C=c|W=w, R=r, D=d)} \quad (5)$$

Estimating $\Pr(I_{rc} = 1)$ using (5) is equivalent to solving the following estimating equation:

$$E\left(\frac{1}{\Pr(R|W)\Pr(C|W, R, D)}(I - m(R, C|\beta))\right)=0 \quad (6)$$

where $m(r, c|\beta)$ denotes a saturated (non-parametric) model for the unknown probability $\Pr(I_{rc} = 1)$.

Here, as in Section 5.1, we estimate the numerator in the right hand side of Equation 5 by the empirical distribution of the data, and the denominator by fitting a logistic regression model for $\Pr(C = c|W = w, R = r, D = d)$. We again use the bootstrap to generate confidence intervals for such estimates.

6. Time-dependent analysis

In this section, we extend the IPTW estimators of Section 5 to provide a time-dependent analysis of direct effects, again using the MIRA trial as an illustration. The description below is an application of the estimation methodology for causal inference based on longitudinal data given in Section 6.4 of van der Laan and Robins (2003). This methodology relies on a marginal structural models for multiple interventions (Hernán et al., 2001; Robins et al., 2003). Note that a single time-dependent regression equation adjusting for causal intermediates would suffer from the problem of confounding by causal intermediates discussed in Section 5.2 at every time point. The analysis presented in this section adequately deals with measured, time-dependent confounders.

6.1. Why We Chose the Inverse Probability of Treatment Weighted (IPTW) Estimator

We started our analysis by implementing the time-dependent inverse probability of treatment weighted (IPTW) estimator for direct effects, defined below. We chose to first implement this estimator, rather than other possible estimators (including G-computation estimators (Robins, 1986), doubly robust estimators (van der Laan and Robins, 2003), and targeted maximum likelihood estimators (van der Laan and Rubin, 2006)). We made this choice because the models required to implement the IPTW estimator, in this application, were easier to specify than the models required by these other methods, and the IPTW is relatively easy to implement. The IPTW estimator requires modelling the probability of condom use at visit v , given observed covariates at or before visit v . In contrast, the G-computation estimator would require models for the probability of HIV infection given past covariates and, at each visit v , models for the joint probability of all time-dependent confounders measured at visit v given past covariates. The doubly robust estimator and targeted maximum likelihood estimator would require models for all of the aforementioned probability distributions. All of the above methods require, in addition, that if a marginal structural model is used it must be correctly specified. We note that the doubly robust estimator and targeted maximum likelihood estimator are generally more efficient and robust to model misspecification than the IPTW estimator; we discuss in Section 8 why we decided, after having computed the IPTW estimator along with confidence intervals for it, not to continue and implement these more sophisticated estimators.

6.2. Observed Data

At enrolment, each subject's baseline covariates W (as listed in Section 5) were measured, prior to study arm assignment. At each of up to 8 follow-up visits, time-dependent confounders (as listed in Section 5), condom use at last sex, and HIV status were recorded; for follow-up visit v (where $v = 0$ refers to enrolment), let $D(v)$ and $C(v)$ denote the values of the time-dependent confounders, and condom use at last sex, respectively. Though in our analyses, $D(v)$ is a vector containing all the measured time-dependent confounders (including diaphragm use), we will refer to it below as “diaphragm use at visit v ” in order to simplify our exposition. Measurement of HIV infection status at visit v is denoted by $I(v)$. Ignoring missing data for the moment, data collected on a single participant from enrolment ($v = 0$) through the last follow-up visit ($v = 8$) is then represented as:

$$O = (W, R, D(0), C(0), I(0), D(1), C(1), I(1), \dots, D(8), C(8), I(8)). \quad (7)$$

Based on this data, the goal is to estimate of the effect of study arm assignment on the probability of HIV infection after two years of participation in the trial, under fixed patterns of condom use over time.

The time-ordering implied by our representation (7) of each participant's data, that baseline confounders W precede randomization R , which precedes decision to use diaphragm at last sex before visit 0, etc., have important consequences. If these variables do not actually occur in this order, our analysis will be biased. As a sensitivity analysis to our choice of time-ordering, we also ran our direct effect analyses using a time-ordering in which $C(i)$ precedes $D(i)$, and got nearly identical results to the analysis using the time-ordering (7). However, in reality the decision to use condoms and/or diaphragms will be more complex; whether such a decision is adequately captured by our time ordering assumption is, unfortunately, not testable from the data collected.

Also, we need to lag all measurements except our HIV status measure by one visit, since HIV infection is detectable by the tests used in the study within a 3 month window, and usually takes at least 3 weeks to be detectable. Though this lag was included in our data analyses, we omit it in explanations of our estimator below for clarity.

We represent the set of observations of condom use over all visits $C(0), C(1), \dots, C(8)$ by the more compact notation $\bar{C}$, and similarly for $\bar{D}$ and $\bar{I}$. We use the notational convention that the first k components of a vector $\bar{A}$ is represented by $\bar{A}(k)$; that is, $\bar{A}(k) = (A(1), \dots, A(k))$. We denote the vector of length 8 of all 0's by $\bar{0}$, and the corresponding vector of all 1's by $\bar{1}$.

6.3. Potential Outcomes

As done for the single time point case in Section 4.2, we define direct effects in terms of potential outcomes. For each possible combination of fixed treatment assignment, r , and condom use over time $\bar{c}$, we let $I_{r\bar{c}}(v)$ represent whether a subject would be HIV-infected by visit v were her study arm and condom use constrained to r and $\bar{c}$, respectively. Similarly, we define $D_{r\bar{c}}(v)$ to represent the potential outcome for diaphragm use that would be reported at visit v , had a subject been assigned to study arm r and had her condom use been constrained to $\bar{c}$.

We note that since for each subject we only measure condom use at last sex 8 times during the study, constraining condom use to $\bar{c}$ means constraining condom use only on these 8 occasions. In what follows, when we refer to setting subjects to use condoms “consistently” or “infrequently”, we really mean setting it just for these 8 occasions; “consistent condom use”

refers to $\bar{c} = (1, 1, 1, 1, 1, 1, 1, 1)$ and “infrequent condom use” refers to $\bar{c} = (0, 0, 0, 0, 0, 0, 0, 0)$. It is assumed that condom use reported at last sex at these 8 visits is generally representative of condom use over the entire study period; in this case, the direct effects relative risks represent the effect of the diaphragm intervention when setting condom use throughout the entire study period to “consistently” or “infrequently”. If this assumption does not hold, for example if condom use at last sex just before each study visit were to systematically differ from condom use at other times, then such an interpretation would not be justified.

We want to estimate the probabilities of HIV infection by the end of the trial, for each study arm, in the following two scenarios: (1) were all subjects constrained to infrequently use condoms at last sex before each quarterly visit and (2) were all subjects constrained to consistently use condoms at last sex before each quarterly visit. In terms of the above notation for potential outcomes, these probabilities of HIV infection are $p_{r\bar{c}} := \Pr(I_{r\bar{c}}(8) = 1)$ for $r \in \{0, 1\}$ and $\bar{c} \in \{\bar{0}, \bar{1}\}$. We can use estimates of these probabilities to compute estimates for direct effects relative risks by taking the appropriate ratios, as described in Section 6.6.

6.4. Main Assumptions: Consistency, No Unmeasured Confounders and ETA

We make the following assumptions, which together imply that our direct effects estimator will be consistent and asymptotically normal. First, the consistency assumption is that at each visit v , (1) observed infection status $I(v)$ equals the potential outcome $I_{r\bar{c}}(v)$ whenever a subject's arm assignment $R = r$ and condom use over time $\bar{C} = \bar{c}$ and (2) observed diaphragm use $D(v)$ equals the potential outcome $D_{r\bar{c}}(v)$ whenever $R = r$ and $\bar{C} = \bar{c}$.

Second, we assume that there are no unmeasured confounders of the relationship between condom use and HIV infection over the visits v . More precisely, we assume at each visit v that condom use at last sex is independent of the set of potential outcomes at all visits, conditioned on all observations made before visit v and also on diaphragm use at visit v :

$$C(v) \{I_{r\bar{c}}, D_{r\bar{c}}\} \mid W, R, \bar{D}(v), \bar{C}(v-1), \bar{I}(v-1). \quad (8)$$

We include diaphragm use at visit v in our set of observations that precede condom use at visit v since we think diaphragm use is determined before condom use, as discussed in Section 6.2 above. (8) is a generalization of the assumptions (1) and (4) from the single time point case.

Third, we make the so-called experimental treatment assignment (ETA) assumption that the probabilities $\Pr(C(v) = 1 \mid W, R, \bar{D}(v), \bar{C}(v-1), \bar{I}(v-1))$ do not equal 0 or 1, almost surely, for any values of $W, R, \bar{D}(v), \bar{C}(v-1), \bar{I}(v-1)$.

Finally, we make a time-ordering assumption for our potential outcomes. This assumption is that for each visit v , the potential outcomes for that visit do not depend on the constraints, in term of condom use, that will be imposed at future visits. Formally, the assumption is that for each visit v , the potential outcomes for that visit ($I_{r\bar{c}}(v)$ and $D_{r\bar{c}}(v)$) only depend on r and $\bar{c}(v)$, and not on $\bar{c}(v+1), \dots, \bar{c}(8)$.

6.5. Model for the Hazard of HIV Infection

Some type of modelling is necessary to avoid extremely high variance in our estimator of the direct effects of treatment assignment fixing condom use at “infrequently.” This is due to there being only 488 subjects who reported not using condoms at last sex at all visits. Assuming a model, as we do below, will lead to some bias due to model misspecification. However, we are able to take advantage of important subject matter knowledge in creating the model; in

particular, our model incorporates the known time lag between when an infection occurs and when it is detectable.

We model the hazard of HIV infection for a participant measured at visit v , would she have been randomized to arm r and been constrained to use condoms according to $\bar{c}(v)$. The model is a function of visit number v , treatment assignment r and condom use history $\bar{c}(v)$.

Model for Hazard of HIV Infection

$$\Pr(I_{r\bar{c}}(v)=1|I_{r\bar{c}}(v-1)=0)=m(v, r, \bar{c}(v)|\beta_0), \text{ for all } v, r, \bar{c}(v), \quad (9)$$

where β_0 is an unknown, time-invariant, population parameter vector. This is an example of a marginal structural model (Robins, 1998).

In our particular choice for a low-dimensional model m , we exploit the following information about the test for HIV infection: HIV infection will, with high probability, be detectable within a 3 month window, and usually takes at least 3 weeks to be detectable. Therefore, (1) condom use measured at the current visit should not have an effect on the outcome of the HIV test at that visit and (2) condom use in the previous 3 months may have a large effect on whether HIV is detected at a given visit, compared to condom use farther in the past. Condom use in the past cannot be completely ignored, since it may impact future behaviour, such as future diaphragm use. We incorporated (1) and (2) into the model by collapsing information about an entire condom use history into just the most recent reported condom use and a single number giving the average condom use before this. We let $H(v, r, \bar{c}(v))$ denote the column vector of basis functions used in the model m in (9); this, and the precise details of the model we used in our analyses is given in our technical report Rosenblum et al. (2007), where we work with a slightly different model than above that incorporates sexual abstinence between visits as well as study sites.

Once we estimate the parameter β_0 , as described below, we can use it to estimate the probabilities $p_{r\bar{c}}$ using the product integral (Andersen et al., 1993) for calculating the cumulative survival probability from discrete hazards:

$$p_{r\bar{c}}=\Pr(I_{r\bar{c}}(8)=1)=1 - \prod_{v=1}^8 \{1 - m(v, r, \bar{c}(v)|\beta_0)\} . \quad (10)$$

6.6. Direct Effects Estimator in the Time-Dependent Setting

In this section, we present the estimating equation for the time-dependent setting that we use to estimate the parameter β_0 of the model in (9), which is then used to estimate direct effects relative risks. The estimating equation uses “stabilized inverse probability of treatment weights (IPTW),” (which we refer to more concisely as “stabilized weights”), which are a generalization of standard inverse probability of treatment weights. In our estimating equation, individual subjects contribute information for all visits except those occurring after HIV seroconversion.

It turns out that solving this estimating equation is equivalent to solving a weighted logistic regression, which can be done with standard statistical software. The details of how to solve this type of estimating equation using weighted logistic regression are given in Section 6.4 of van der Laan and Robins (2003).

Before giving the estimating equation used to estimate β_0 , we define the stabilizing weights that will be used in this estimating equation:

$$sw(W, R, \bar{D}(v), \bar{C}(v), \bar{I}(v-1)) = \frac{\prod_{v' \leq v} \Pr(C(v') | R, \bar{C}(v' - 1), I(v' - 1) = 0)}{\prod_{v' \leq v} \Pr(C(v') | W, R, \bar{D}(v'), \bar{C}(v' - 1), \bar{I}(v' - 1))}. \quad (11)$$

They are called “stabilized” since the hope is that for large values of the standard IPTW weights, which correspond to small values in the denominator, the numerator terms will also be small; in this case, one would prevent single observations from having a disproportionate amount of weight in the estimation procedure.

Just as in Sections 5.1 and 5.2, we use logistic regression models to estimate the IPTW weights. We estimate each term in the product in the numerator of (11) by fitting a separate logistic regression model; we do the same for each term in the product in the denominator. We give the precise details of these logistic regression models in our technical report Rosenblum et al. (2007). In order to reduce the variance of the IPTW estimator, we decided prior to looking at the data to truncate all estimated stabilized weights at 25. As we discuss in the Section 7 below, the largest value of any estimated stabilized weight in our analysis was 16.9, so no weights were actually truncated.

Let n denote the number of subjects in the study. Recall that $H(v, r, \bar{c}(v))$ denotes the column vector of basis functions used in the model m in (9). The estimating equation corresponding to the IPTW estimator with stabilized weights is given by

$$\sum_{i=1}^n \sum_v^{l_i} sw(W, R, \bar{D}(v), \bar{C}(v), \bar{I}(v-1)) \{I(v) - m(v, R, \bar{C}(v) | \beta_0)\} H(v, R, \bar{C}(v)) = 0,$$

where the inner sum for each subject i is taken starting at follow-up visit $v = 1$ up to and including the first visit at which the participant tests positive for HIV, denoted by l_i ; for those women who remain HIV negative throughout, the sum is over all eight follow-up visits; that is, $l_i = 8$ for these women. Note the similarity in form to the simpler estimating equations (3) and (6). Details and proof that this estimating equation indeed has mean 0 at the true β under the assumptions given in Sections 6.2-6.5 are given in Section 6.4 of van der Laan and Robins (2003).

Once we obtain a solution $\hat{\beta}$ to the above estimating equation, we can use it to generate estimates $\hat{p}_{r\bar{c}}$ for the probability of HIV infection by the end of the trial for a randomly chosen subject in the study, would she have been assigned to study arm r and have had condom use profile set to $\bar{c}$. Following (10) above, we have

$$\hat{p}_{r\bar{c}} = 1 - \prod_{v=1}^8 (1 - m(v, r, \bar{c}(v) | \hat{\beta}))$$

We can then construct estimates of direct effects relative risks by taking ratios of the estimates $\hat{p}_{r\bar{c}}$. The direct effects relative risk estimate with condom use fixed at “infrequently” is $\hat{p}_{1\bar{0}} / \hat{p}_{0\bar{0}}$. (Recall that $\bar{0}$ denotes the vector of length 8 of all 0's.) The direct direct effects relative

risk estimate with condom use fixed at “consistently” is $\hat{p}_{1\bar{1}}/\hat{p}_{0\bar{1}}$. We used the bootstrap BCa method for constructing confidence intervals for these direct effects relative risk estimates, where replicates were selected using subjects (not subject-visits) as the experimental unit.

7. Results of Our Data Analyses for the MIRA Trial

In this section, we give summary statistics for the MIRA trial and then the results of the time-dependent analysis. Due to space constraints, we don't give the results of the single time point analysis, but we do discuss below how they are qualitatively similar to the results for the time-dependent case.

Table 1 below gives the distribution of the two arms of the study with regard to the number of seroconversions to HIV during follow-up. The intention-to-treat relative risk of HIV infection by the end of the trial is $(158/2472)/(151/2476) = 1.05$ (95% CI: 0.84, 1.30). 6% of participants in the diaphragm arm and 5% of the participants in the control arm were lost to follow-up.

The time-dependent direct effects analysis was based on a discrete time scale made up of the eight quarterly visits by each participant, in addition to an initial enrolment visit (visit 0). For individuals who seroconverted, the (discrete) time to seroconversion was defined as the visit number at the time of the first positive HIV test. For participants who missed one or more visits between the last negative and first positive test, the seroconversion time was defined as the visit number covering the midpoint between the two discordant tests. Covariates were imputed by carrying forward the most recently observed value; for example, if a subject came to visit 3 but missed visits 4 and 5, her covariates for visits 4 and 5 were imputed to be what was recorded at visit 3.

Before calculating the direct effects relative risks of study arm assignment on HIV infection with condom use fixed, we looked at whether the MIRA trial data were consistent with reported condom use being causally related to HIV infections. If it were not, then a direct effects analysis fixing reported condom use would probably not be of much use. We did a causal analysis of the effect of condom use at last sex on subsequent HIV seroconversion, within each study arm separately. We did this by estimating $p_{r\bar{c}}$ for $r \in \{0, 1\}$, $\bar{c} \in \{\bar{0}, \bar{1}\}$ as described in Section 6. But instead of looking at the direct effects relative risks holding condom use fixed as described in that section, we looked at the following relative risk: $p_{r\bar{c}}/p_{r\bar{c}'}$, first for $r = 0$, $\bar{c} = \bar{1}$, $\bar{c}' = \bar{0}$, which is the relative risk of HIV infection comparing the scenarios in which subjects are constrained to use condoms consistently vs. infrequently, in the control arm. We then calculated this relative risk $p_{r\bar{c}}/p_{r\bar{c}'}$, except setting $r = 1$, $\bar{c} = \bar{1}$, $\bar{c}' = \bar{0}$, which gives the relative risk for the diaphragm arm. The estimates for these measures of condom efficacy are: 0.33 (95% CI 0.16, 2.33) for the control arm and 0.54 (95% CI: 0.30, 0.98) for the diaphragm arm, indicating a statistically significant protective effect in the diaphragm arm but no statistically significant result in the control arm. This protection against HIV infection is less than the approximately 80% relative protection normally ascribed to condom use (Weller and Davis-Beaty, 2002). The lower protection reflected in the above estimates may be due to overreporting of condom use due to social desirability bias, which would dilute the measured effectiveness of condoms.

We now present the results of the time-dependent direct effects analysis described in Section 6. Using the methods from Section 6, we can generate an estimate of $p_{r\bar{c}} := \Pr(I_{r\bar{c}} = 1)$ for $r = 0, 1$ and any specific pattern of condom use reported over the 8 possible visits, which is denoted by $\bar{c}$. For example, $\bar{c} = (1, 1, 1, 1, 1, 1, 1, 1)$ represents condom use at last sex at all 8 visits. Here we focus on two extreme patterns of condom use: (i) condoms used at last sex at all visits, so that $c(v) = 1$ for all v , and (ii) condoms not used at last sex at all visits, so that $c(v) = 0$ for all v .

Table 2 gives the results for these two extreme patterns of condom use. The first row of Table 2 gives the estimated direct effects of assignment to the diaphragm arm vs. control arm, for condom use set to “no” at all 8 visits. First, the direct effects relative risk 0.59(95% CI: 0.23, 3.17) is given; this is an estimate of $p_{1\bar{c}}/p_{0\bar{c}}$ for $\bar{c} = \bar{0}$. We also give the direct effects cumulative survival difference 0.10(95% CI: -0.13, 0.33), which is an estimate of $(1 - p_{1\bar{c}}) - (1 - p_{0\bar{c}})$ for $\bar{c} = \bar{0}$. The second row of Table 2 gives the estimated direct effect of assignment to the diaphragm arm vs. control arm, for condom use set to “yes” at all 8 visits. All confidence intervals were calculated with the non-parametric bootstrap BCa method, using 10,000 iterations. These results are qualitatively similar to the results for the single time point analysis (given in our technical report (Rosenblum et al., 2007)), in that none of them are statistically significant, and all estimates have very wide confidence intervals. For comparison to the direct effects estimates, recall that the intention-to-treat relative risk of HIV infection by the end of the trial is $158/151 = 1.05$ (95% CI: 0.84, 1.30).

Though both point estimates are in the direction of assignment to the diaphragm arm being protective in these scenarios, neither is close to statistical significance. Also, based on the bootstrap estimates of standard errors, both of the direct effects relative risk point estimates are within one standard error of the value corresponding to no effect of assignment to the diaphragm arm. We therefore conclude that the direct effects analysis provides no evidence in support of, nor in refutation of, diaphragms providing a protective effect against HIV infection.

We briefly discuss the stabilized weights, and how these may have impacted the widths of our confidence intervals. The estimated, stabilized weights in our analysis were distributed as follows: Min: 0.000014, 1st Quartile: 0.42, Median: 0.77, Mean: 0.89, 3rd Quartile: 1.14, Max: 16.89. Since none of the weights exceeded 25, none were truncated. As a sensitivity analysis to determine if the large confidence intervals in Table 2 were a result of our stabilized weights, we ran our time-dependent analysis with no weights at all (that is, all weights = 1). The confidence intervals were still quite large; for example, the 95% confidence intervals for the direct effects relative risks were (0.348, 2.12) for condom use set to “infrequently” and (0.55, 1.18) for condom use set to “consistently.” We conclude that the large confidence intervals are not due primarily to the stabilized weights we used. We note that a more refined analysis could include data-adaptive selection of the truncation threshold, so as to optimally tradeoff bias and variance (see e.g. Bembom and van der Laan (2008)). Since no weights were extremely large, and based on this sensitivity analysis, we decided not to implement this refinement.

8. Discussion

The direct effects estimator we used, based on the inverse probability of treatment weighted (IPTW) estimator, could be biased if (1) there are unmeasured confounders (e.g. characteristics of male partners associated with their condom use and HIV status); (2) there is measurement error in reported condom use (for example, due to social desirability bias, or if quarterly reported condom use at last sex is not sufficiently informative about overall condom use) or measurement error in confounders; (3) the models for condom use or hazard of HIV infection are not correctly specified; (4) missing data values have a different distribution than observed values; (5) the experimental treatment assignment assumption is violated (see Section 6.4); or (6) the consistency assumption or time-ordering assumption is violated. All of these biases may exist to some degree, with (1) and (2) being of particular concern in the context of the MIRA trial. (1) is important since partners, who are believed to control condom use much of the time, may be associated with important unmeasured confounding. (2) is important since condom use by self-report is likely to be inflated upward, and this is consistent with our causal analysis of condom use in Section 7. We treat the point estimates for direct effects relative risks with a good deal of scepticism for these reasons.

However, the wide confidence intervals in the direct effects analysis allow us to make the following valuable conclusion: The data in the MIRA trial do not contain enough information to determine whether differential condom use may have masked a possibly protective effect of study arm assignment on HIV infection. This was not at all evident from a first look at the intention-to-treat results and the large difference in reported condom use, which on their face seemed, to the contrary, to imply that there might be enough information in the MIRA trial data to adjust for condom use and show a positive effect of diaphragms. The crude analysis given in Section 2 to adjust for condom use also suggested a possible protective effect of the diaphragm intervention. Such a crude analysis did not account for confounding or the time-dependent nature of the data; an analysis taking these into account (such as our direct effects analysis) was required to determine if there was enough information in the MIRA trial data to adjust for condom use and possibly uncover a positive effect of diaphragms. Without such an analysis, readers of the Lancet article reporting only the MIRA trial intention-to-treat results and differential condom use (Padian et al., 2007) may continue to wonder whether there was evidence in the trial data for a benefit of the diaphragm intervention.

The question remains as to how to reconcile the finding from our direct effects analysis with the result of our crude analysis from Section 2, which suggested an effect of the diaphragm intervention when controlling for condom use. Recall in the crude analysis from Section 2 that we assumed consistent condom use reduces HIV risk by 80%, in accordance with prior studies of condom effectiveness (Weller and Davis-Beaty, 2002). However, our point estimate of the causal effect of setting condom use to consistent vs. infrequent from Section 7 above corresponds to *reported* consistent condom use reducing HIV risk by $(100\% - 54\%) = 46\%$; that this value is much lower than 80% can be explained by bias in self-reporting of condom use. Taking this information (that was obtained only after our causal analysis) into account, if we substitute 46% instead of 80% protection due to consistent condom use into the crude analysis in Section 2, we get that increasing condom use in the diaphragm arm to the level observed in the control arm would lead to approximately ≈ 23 infections prevented. However, this is much lower than the number of prevented infections required to achieve a statistically significant difference between the two study arms, at significance level $\alpha = 0.05$ (two-tailed); 39 such infections would be required to get a statistically significant result. Thus, once we take account of reporting bias by using the information from the direct effects analysis regarding the causal link between condom use and HIV infection, we get that the crude analysis from Section 2 agrees with the direct effects analysis.

In Section 6.1, we explained why we chose to start our time-dependent analysis by implementing the IPTW estimator for direct effects, rather than other possible estimators. Had the IPTW-based analysis provided evidence that the diaphragm intervention had an effect on HIV outcome at fixed levels of condom use, the next step would have been to implement a more efficient, robust estimator, such as the doubly-robust estimator (van der Laan and Robins, 2003) or the targeted maximum likelihood estimator (van der Laan and Rubin, 2006). These estimators are robust in that they give asymptotically consistent estimates whenever either (1) correctly specified models are used for the probability of condom use at visit v given observed covariates at or before visit v , or (2) correctly specified models are used for the probability of HIV infection given past covariates and, at each visit v , models for the joint probability of all time-dependent confounders measured at visit v given past covariates. Implementation of either of these more sophisticated estimators for longitudinal data sets requires considerably more work than implementing the IPTW estimator; generally the gain in terms of robustness to model misspecification and in efficiency from these estimators more than outweighs the extra work required to implement them. In this application, however, the IPTW analysis showed that there was probably little, if any, signal in the data that would indicate a causal effect of the diaphragm intervention on HIV outcome; in particular, both point estimates for the direct effects relative risks were within one standard error of 1. This lack of signal in the data regarding an effect of

the diaphragm intervention is corroborated by the intention-to-treat analysis, as well as several other analyses considering different endpoints thought to be indicative of diaphragm effectiveness, all of which failed to show an effect. We decided not to proceed to implement a more efficient, robust estimator, such as the targeted maximum likelihood estimator, since we judged it highly unlikely to show an effect using this data. For the same reason, we decided not to implement a less biased but more involved method of imputing missing data values, such as multiple imputation or modeling the censoring mechanism.

We briefly discuss some alternative analyses that could be done to answer some important public health questions related to the MIRA trial. For example, one could try to estimate diaphragm efficacy at a fixed level of condom use, that is, the effect of diaphragm use D on HIV status I , fixing condom use C . We expect this would differ from what we estimated using direct effects, which involved the effect of being assigned to the diaphragm arm. An analysis of the effect of diaphragm use D on I , fixing condom use C would require additional assumptions than our direct effects analysis, since one would have to deal with time-dependent confounding of the effect of diaphragm use on HIV status. This highlights an advantage of doing a direct effects analysis of treatment assignment in a randomized trial, as we did in this paper: there cannot be confounders of (randomized) treatment assignment and primary outcome. Another possible analysis would involve estimating natural direct effects (Robins and Greenland, 1992), rather than the controlled direct effects estimated in this paper. This would correspond to estimating the relative risk of HIV infection between the two study arms, were all subjects in the diaphragm arm to use condoms with the same frequency as they would have if they had instead been assigned to the control arm. We think this would be an interesting parameter, since it would shed light on the effectiveness of the diaphragm intervention with condom use “fixed” in a different sense than in controlled direct effects; we note, however, that estimating natural direct effects requires stronger assumptions than for estimating controlled direct effects (Petersen et al., 2006).

Another type of causal analysis that could be applied to the MIRA trial is principal stratification (Frangakis and Rubin, 2002). It could be used, for example to estimate the effect of study arm assignment for just the group of women who would not have used condoms regardless of which study arm they had been assigned to. This parameter may shed light on the important public health question of whether a diaphragm intervention would provide protection against HIV for women who cannot get their partners to consistently use male condoms, or sustain high levels of condom use over time. A different set of assumptions than for the direct effects analysis would be necessary to estimate this parameter (Frangakis and Rubin, 2002); also, R cannot be assumed to be an instrumental variable, since R potentially influences I through the separate causal pathway involving C .

The direct effects analysis can also be applied in trials with blinding, to try to remove distortion caused by a secondary intervention given to both arms. For example, consider microbicide trials in which a placebo (an inactive gel) is used, and participants cannot distinguish whether they are using the active microbicide or placebo, as discussed in Mantell et al. (2005) and Trussell and Dominik (2005). In a truly blinded trial, assignment to different study arms cannot affect the outcome through systematically modifying subjects' risk behaviours (such as condom use). But even in blinded trials, secondary interventions can affect the intention-to-treat estimator, pushing it either toward or away from the efficacy of the study product, as shown in Trussell and Dominik (2005). A direct effects analysis can remove some of this distortion caused by a secondary intervention, if this secondary intervention influences the outcome only through a measured, potentially controllable behaviour (e.g. condom use). A direct effects analysis produces an estimate of what the effect of treatment assignment would have been, had subjects been constrained in this behaviour. This information could be useful, for example, in a blinded microbicide trial with a strongly positive intention-to-treat result; here, a direct effects

analysis could try to uncover whether intervention effectiveness might be very different in a setting without condom counselling in which condom use were quite low. Similarly, in a blinded microbicide trial that resulted in a null intention-to-treat result, the direct effects analysis could elucidate whether condom counselling could have pushed the intention-to-treat result far away from actual microbicide efficacy.

Changes to the MIRA trial design could help alleviate some of the difficulties discussed in this paper. For example, more resources could be put into recruiting participants whose partners refuse to use condoms. This is the population for which the diaphragm intervention was initially conceived, but in practice is difficult to recruit. Such a population would have low condom use in both arms in the presence of intensive condom counselling, obviating some of the problems discussed in this paper. To identify such subjects, recruiters could screen for women having covariates that in the MIRA trial were highly predictive of low condom use. There is much need for trial designs that are better able to target efficacy and effectiveness of new HIV prevention methods.

Acknowledgments

We would like to thank all the women who participated in the MIRA trial, the entire study team and the collaborating institutions, including staff and investigators at the University of Zimbabwe-University of California Collaborative Research Programme (Dr. T. Nhemechena and Dr. T. Chipato), at the Perinatal HIV Research Unit of the University of Witwatersrand (Dr. G. De Bruyn, Dr. G. Gray and Dr. J. McIntyre), at the South African Medical Research Council, HIV Prevention Research Unit (Dr. G. Ramjee), at Ibis Reproductive Health (Ms. K. Blanchard) and staff at the University of California San Francisco. Special thanks to Helen Cheng for statistical support. The MIRA trial was funded through a grant from the Bill and Melinda Gates Foundation (#21082).

Michael Rosenblum was supported by a Ruth L. Kirschstein National Research Service Award (NRSA) under NIH/NIMH grant 5 T32 MH-19105-19. This work was supported in part by Grant #AI070-043 from the National Institute of Allergy and Infectious Diseases, Bethesda, Maryland, USA. The second author also gratefully acknowledges the support of the Rockefeller Foundation through a Fellowship at the Study Center in Bellagio during the development of this work. Mark van der Laan was supported by NIH grant R01 A1074345-01.

References

- Andersen, P.; Borgan, O.; Gill, R.; Keiding, N. Statistical Models Based on Counting Processes. NY, USA: Springer; 1993.
- Anderson J, Rietmeijer C, Wilson R. Asking about condom use: is there a standard approach that should be adopted across surveys? Annual meeting of the American Statistical Association, St Louis. 1998
- Angrist JD, Imbens GW, Rubin DB. Identification of causal effects using instrumental variables. J Am Statist Ass 1996 June;91(434):444–455.
- Bembom, O.; van der Laan, MJ. Data-adaptive selection of the truncation level for inverse-probability-of-treatment-weighted estimators. UC Berkeley Division of Biostatistics Working Paper Series. Working Paper 230. 2008. <http://www.bepress.com/ucbbiostat/paper230>
- Dawid AP. Causal inference without counterfactuals. J Am Statist Ass 2000 June;95(450):407–424.
- Efron, B.; Tibshirani, R. An Introduction to the Bootstrap. NY, USA: Chapman and Hall; 1993.
- Feinstein, AR. The architecture of clinical research. Philadelphia: W.B. Saunders; 1985.
- Frangakis C, Rubin D. Principal stratification in causal inference. Biometrics 2002;58:21–29. [PubMed: 11890317]
- Frangakis CE, Rubin DB. Addressing complications of intention-to-treat analysis in the combined presence of all-or-none treatment-noncompliance and subsequent missing outcomes. Biometrika 1999;86(2):365–379.
- Friedman, L.; Furberg, C.; DeMets, D. Fundamentals of Clinical Trials. 3rd. NY, USA: Springer; 1998.
- Hernán M, Brumback B, Robins J. Marginal structural models to estimate the causal effect of zidovudine on the survival of hiv-positive men. Epidemiology 2000;11(5):561–570. [PubMed: 10955409]

- Hernán M, Brumback B, Robins JM. Marginal structural models to estimate the joint causal effect of nonrandomized treatments. *Journal of the American Statistical Association – Applications and Case Studies* 2001;96(454):440–448.
- Hirano K, Imbens GW, Rubin D, Zhou XH. Assessing the effect of an influenza vaccine in an encouragement design. *Biostatistics* 2000;1(1):69–88. [PubMed: 12933526]
- Jewell, NP. *Statistics for Epidemiology*. Boca Raton, Florida, USA: Chapman and Hall/CRC; 2004.
- Mantell J, Myer L, CD A, et al. Microbicide acceptability research: current approaches and future directions. *Soc Sci Med* 2005;60:319–330. [PubMed: 15522488]
- Padian N, Straten A van der, Ramjee G, Chipato T, Bruyn D de, Blanchard K, Shiboski S, Montgomery E, Fancher H, Cheng H, Rosenblum M, Laan M van der, Jewell N, McIntyre J, Mira team. Diaphragm and lubricant gel for prevention of hiv acquisition in southern african women: a randomised controlled trial. *The Lancet* 2007 July;370(9583):251–261.
- Pearl, J. *Causality: Models, Reasoning, and Inference*. Cambridge, UK: Cambridge University Press; 2000a.
- Pearl, J. Direct and indirect effects. In: Kaufmann, M., editor. *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*. 2000b. p. 411–420.
- Petersen M, Sinisi S, van der Laan M. Estimation of direct causal effects. *Epidemiology* 2006;17:276–284. [PubMed: 16617276]
- Robins J. A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect. (with errata). *Mathematical Modelling* 1986;7:1393–1512.
- Robins J. Marginal structural models. 1997 *Proceedings of the American Statistical Association* 1998:1–10.
- Robins J, Finkelstein D. Correcting for non-compliance and dependent censoring in an aids clinical trial with inverse probability of censoring weighted (ipcw) log-rank tests. *Biometrics* 2000;56(3):779–788. [PubMed: 10985216]
- Robins J, Greenland S. Identifiability and exchangeability for direct and indirect effects. *Epidemiology* 1992;3:143–155. [PubMed: 1576220]
- Robins, JM. *Causal Inference from Complex Longitudinal Data*. *Latent Variable Modeling and Applications to Causality Lecture Notes in Statistics* (120). Berkane, M., editor. NY: Springer Verlag; 1997.
- Robins JM, Hernán M, Siebert U. Effects of multiple interventions. *Population Health Metrics* 2003;2:2191–2230.
- Rosenbaum P. The consequences of adjustment for a concomitant variable that has been affected by the treatment. *JR Statist Soc A* 1984;147(5):656–666.
- Rosenblum, M.; Jewell, N.; van der Laan, M.; Shiboski, S.; van der Straten, A.; Padian, N. Detailed version: Analyzing direct effects of treatments in randomized trials with secondary interventions: An application to hiv prevention trials. UC Berkeley Division of Biostatistics Working Paper Series. 2007 Oct. <http://www.bepress.com/ucbbiostat/paper225> (Working Paper 225)
- Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 1974;66(2):688–701.
- Sheiner LB, Rubin DB. Intention-to-treat analysis and the goals of clinical trials. *Clin Pharmacol Ther* 1995;57:6–15. [PubMed: 7828382]
- Trussell J, Dominik R. Will microbicide trials yield unbiased estimates of microbicide efficacy? *Contraception* 2005;72:408–413. [PubMed: 16307961]
- van der Laan, M.; Robins, M. *Unified Methods for Censored Longitudinal Data and Causality*. Springer-Verlag; 2003.
- van der Laan, M.; Rubin, D. Targeted maximum likelihood learning. *The International Journal of Biostatistics*. 2006. (Article 11) Available at: <http://www.bepress.com/ijb/vol2/iss1/11>
- Weller SC, Davis-Beatty K. Condom effectiveness in reducing heterosexual hiv transmission. art. no.: Cd003255. *Cochrane Database of Systematic Reviews* 2002;110.1002/14651858.cd003255

Figure 1.
Direct Effects Causal Diagram

Figure 2.
Direct Effects Causal Diagram with Pre-Treatment Confounder W

Figure 3.
Direct Effects Diagram with Confounder as Causal Intermediate

Figure 4.
Causal Diagram with Confounder as Causal Intermediate and Unobserved Variable U.

Table 1

Summary of HIV Seroconversions by Study Arm

	Number of Participants	Seroconversions (percent of group)	Woman-years of Follow-up	Incidence Rate (per 100 woman-years)
Total	4948	309 (6.2%)	7702	4.0
Intervention	2472	158 (6.4%)	3835	4.1
Control	2476	151 (6.1%)	3867	3.9

Table 2

Estimates of the Direct Effect of Assignment to Study Arm on HIV Infection Using a Time-Dependent Analysis

	Relative Risk by Visit 8	95% CI	Cum. Survival Difference	95% CI
Condom Use Set to “No” all Visits	0.59	(0.23, 3.17)	0.10	(-0.10, 0.39)
Condom Use Set to “Yes” all Visits	0.96	(0.61, 1.48)	0.00	(-0.03, 0.03)