Published in final edited form as:

J R Stat Soc Series B Stat Methodol. 2017 June ; 79(3): 719–735. doi:10.1111/rssb.12188.

Causal analysis of ordinal treatments and binary outcomes under truncation by death

Linbo Wang^{*,1}, Thomas S. Richardson¹, and Xiao-Hua Zhou^{1,2}

¹University of Washington, Seattle, USA

²Veterans Affairs Puget Sound Health Care System, Seattle, USA

Abstract

It is common that in multi-arm randomized trials, the outcome of interest is “truncated by death,” meaning that it is only observed or well-defined conditioning on an intermediate outcome. In this case, in addition to pairwise contrasts, the joint inference for all treatment arms is also of interest. Under a monotonicity assumption we present methods for both pairwise and joint causal analyses of ordinal treatments and binary outcomes in presence of truncation by death. We illustrate via examples the appropriateness of our assumptions in different scientific contexts.

Keywords

Bayesian analysis; Causal inference; Multi-arm trials; Ordinal treatment variable; Principal stratification; Survey incentives

1 Introduction

In multi-arm randomized trials, researchers are often interested in analyzing treatment effects on an outcome that is measured or well-defined only when an intermediate outcome takes certain values (Robins, 1986; Rubin, 2000, 2006; Egleston et al., 2007; Chiba and VanderWeele, 2011; Ding et al., 2011). For example, consider a multi-arm randomized HIV vaccine trial. Scientists might be interested in evaluating vaccine effects on HIV viral load as it correlates with infectiousness and disease progression (Hudgens et al., 2003; Gilbert et al., 2003). However, HIV viral load is typically measured only for infected individuals. Two problems occur in this case: in general, there are many potential comparisons that can be made between different vaccination groups among infected subjects; moreover, these comparisons are subject to selection bias as the vaccine may affect susceptibility to HIV infection. In the simple case of a two-arm trial, to deal with the selection bias problem, several authors have proposed to consider the vaccine effects on viral load among the always-infected stratum, the subpopulation who would become infected regardless of whether they are vaccinated or not (e.g., Hudgens et al., 2003; Gilbert et al., 2003). However, there has not been much work on analyzing this type of trial with more than two arms.

*Address for correspondence: Linbo Wang, Department of Biostatistics, University of Washington, Box 357232, Seattle, WA 98195, USA lbwang@uw.edu.

By convention, the intermediate outcome is called “survival,” and we say the final outcome is truncated by death if it is only observed and/or well-defined for “survivors.” Thus in the HIV vaccine example above, the always-infected stratum is referred to as the “always-survivor” stratum. The causal contrast among the always-infected subjects is hence called the (always-)survivor average causal effect (SACE) (Rubin, 2000).

In general, even in a two-arm trial, the SACE is not identifiable without strong untestable assumptions. As a result, there are no consistent tests for detecting non-null vaccine effects in the always-infected stratum. Instead, under some reasonable assumptions, Hudgens et al. (2003) tested the null hypothesis presuming the maximal degree of selection bias. Their approach is related to estimation of bounds on SACE, which has been extensively studied in literature. For example, Zhang and Rubin (2003) developed bounds on SACE under various assumptions including the monotonicity assumption and the stochastic dominance assumption. Imai (2008) provided an alternative proof that the bounds of Zhang and Rubin (2003) are sharp by formulating the truncation-by-death problem as a “contaminated data” problem. These testing and estimation methods are appealing in practice as they don’t rely on strong identifiability assumptions.

However, to the best of our knowledge, there has not been much discussion on testing and estimation of SACEs in a multi-arm trial, which is fairly common in medical practice (Schulz and Grimes, 2005). Prior to our work, Lee et al. (2010) considered a sensitivity analysis approach to identify all SACEs in a three-arm trial. Their identification results rely on a strong parametric assumption and several sensitivity parameters. In this article, we instead propose a framework to systematically analyze SACEs in a general multi-arm trial without strong identification assumptions. To the best of our knowledge, our method is also the first that is readily applicable to randomized trials with more than three treatment arms under truncation by death.

The testing and estimation of SACEs in a multi-arm trial are more challenging compared to two-arm trials. Firstly, in general there are many different SACEs that are well-defined. As we show later in Section 2.2, consideration of all SACEs (as in Lee et al. (2010)) can lead to paradoxical non-transitive conclusions. Hence we instead restrict our attention to comparisons within the “finest” (principal) strata, thereby avoiding this difficulty. Secondly, one needs to distinguish between an overall analysis of treatment effects and a separate analysis of each individual contrast. In the simple setting without truncation by death, it is widely known that compared to all pairwise comparisons with correction for multiple comparisons, an overall analysis such as an ANOVA test often provides more power for testing the overall treatment effect in a multi-arm trial. When truncation by death is present, due to non-identifiability of SACEs, this advantage becomes more fundamental as it remains even when the sample size goes to infinity. In contrast to Lee et al. (2010), we distinguish between simultaneous versus marginal inference for SACEs, and argue that they should be used to answer different questions. In particular, we show that compared to marginal inference procedures, our proposed simultaneous inference procedures provide more power for testing the overall treatment effect and the advantage remains even with an infinite sample size. Thirdly, the simultaneous inference problem is unique to a multi-arm trial. Again, since SACEs are not identifiable, traditional statistical inference tools for multi-arm

trials without truncation by death are not directly applicable to our setting. Instead, we develop novel simultaneous inference procedures to test an overall treatment effect, and show that they have desirable asymptotic properties. We also generalize the marginal inference procedures for a two-arm trial to get sharp bounds on SACEs for a general multi-arm trial. To focus on addressing these challenges, in this paper, we restrict our attention to trials with ordinal treatment groups and binary outcomes.

The rest of this article is organized as follows. In Section 2, we introduce our notations, assumptions and define our causal estimands. We also address the transitivity issue and identify three specific testing and estimation questions that may arise in a general multi-arm trial with truncation by death. We then propose three novel procedures that answer these questions in Section 3, 4 and 5, respectively. In Section 3, we discuss the unique challenges for hypothesis testing with non-identifiable parameters, and develop a novel step-down testing procedure to test the overall treatment effect in this situation. In Section 4, we develop a linear programming algorithm to test an overall clinically relevant treatment effect. In Section 5, we derive the sharp marginal bounds for each causal contrast of interest. In Section 6, we illustrate the proposed procedure with real data analyses. Results from simulation studies can be found in the the Supplementary Materials. We end with a discussion in Section 7.

2 Framework

2.1 Data structure and assumptions

Consider a multi-arm trial with control and multiple arms of active treatment. Let Z be an ordinal treatment variable, where $Z = 0$ corresponds to the control treatment, and $Z \in \{1, \dots, m\}$ corresponds to different arms of active treatment. In what follows, we use the terminology “treatment arms” and “treatment levels” interchangeably. We assume that each subject has $m + 1$ dichotomous potential outcomes $Y(z)$, $z = 0, \dots, m$, where $Y(z)$ is defined as the outcome that would have been observed if the subject was assigned to treatment arm z . Similarly, we define $S(z)$ as the potential survival status under treatment assignment z . We assume $Y(z)$ is well-defined only if $S(z) = 1$. In other words, the outcome of interest is well-defined only for subjects who survive to the follow-up visit. We also assume that the observed data $(Z_i, S_i, Y_i; i = 1, \dots, N)$ are independently drawn from an infinite super-population.

Let $G = (S(0), \dots, S(m))$ denotes the *basic principal stratum* (Frangakis and Rubin, 2002). If we let the letter L denote $S(z) = 1$ (meaning “live”) and the letter D denote $S(z) = 0$ (meaning “die”), then G can be rewritten as a string consisting of the letters “L” and “D.” For example, in a three-arm trial, $G_i = DLL$ indicates that subject i would die under control, but would survive under active treatment 1 or 2.

We make the following assumptions.

Assumption 1—Stable unit treatment value assumption (SUTVA (Rubin, 1980)): there is no interference between units, and there is only one version of treatment.

Under the SUTVA, the observed outcome equals the potential outcome under the observed treatment arm, namely $Y = Y(Z)$ and $S = S(Z)$.

Assumption 2—Random treatment assignment: $Z \perp (S(0), \dots, S(m), Y(0), \dots, Y(m))$.

Assumption 3—Monotonicity: $S_i(z_1) \leq S_i(z_2)$, $i = 1, \dots, N$, $z_1 \leq z_2$.

The monotonicity assumption is sometimes plausible in social science studies if the treatment options can be reasonably ordered. For example, in randomized experiments evaluating the effect of incentives on survey response quality, it is intuitive that higher level of incentives would not hurt survey response rates. This assumption tends to be more controversial in medical studies, in which there are often trade-offs between treatment benefits and side effects.

The only possible strata under the monotonicity assumption are strata of the form $D \cdots DL \cdots L$. To compress notation, we denote all possible principal strata as $(D^k L^{m+1-k}; k = 0, \dots, m+1)$, where members of principal stratum $D^k L^{m+1-k}$ would die if assigned to the first k treatment arms but would survive if assigned to the remaining $m+1-k$ treatment arms.

2.2 Causal estimands and questions

For randomized trials with two treatment arms, it is common to estimate the average causal effect in the LL stratum (Kalbfleisch and Prentice, 1980; Robins, 1986; Rubin, 2000), the only subgroup for which both of the potential outcomes are well-defined: $SACE = E[Y(1) - Y(0) | G = LL]$. In a general multi-arm trial, researchers may be interested in comparisons of potential outcomes within the same basic principal stratum. For example, in the case where we have three levels of treatment: 0, 1, 2, the target estimands are $E[Y(2) - Y(1) | G = LLL]$, $E[Y(1) - Y(0) | G = LLL]$, $E[Y(2) - Y(0) | G = LLL]$ and $E[Y(2) - Y(1) | G = DLL]$. These contrasts are causally meaningful as the memberships of basic principal strata are defined at baseline.

To define the causal estimands for a general multi-arm trial, we first introduce some notation. Let $\mu_g^z \equiv E[Y(z) | G=g]$ denote the mean potential outcome under treatment assignment z in basic principal stratum g . Also, let $\mathcal{M}(g)$ denote the minimal treatment level under which members of principal stratum g can survive. In other words, for members of principal stratum g , $S(z) = 1$ if and only if $z \geq \mathcal{M}(g)$. Consequently, μ_g^z is well-defined if and only if $z \geq \mathcal{M}(g)$. Under the monotonicity assumption, all basic principal strata take the form $g = D^k L^{m+1-k}$. By definition, $\mathcal{M}(D^k L^{m+1-k}) = k$. Also let $\Omega_k = \{g: \mathcal{M}(g) \leq k\}$ denote the collection of basic principal strata whose members would survive if assigned to treatment arm k . The pairwise causal estimands in a multi-arm trial then take the form

$$\Delta(z_1, z_2; g) \equiv \mu_g^{z_1} - \mu_g^{z_2}, \text{ where } g \in \Omega_{m-1}, z_1 > z_2 \geq \mathcal{M}(g). \quad (1)$$

For notational simplicity, in this article, when we write the notation μ_g^z and $(z_1, z_2; g)$, we always assume that it is well-defined. We also note that the parameters involved in defining

the causal contrasts $(z_1, z_2; g)$ are contained in the parameter vector

$$\mu_{m-1} \equiv \left(\mu_g^z; g \in \Omega_{m-1}, z \geq \mathcal{M}(g) \right).$$

There are other meaningful causal contrasts that are made within *coarsened principal strata*, defined as groups that combine several basic principal strata (Cheng and Small, 2006). For example, in the case of a three-arm trial, the contrast $E[Y(2) - Y(1) | G \in \{LLL, DLL\}]$ is also causally meaningful as memberships of the coarsened principal strata $\{LLL, DLL\}$ are also defined at baseline. Some previous researchers hence consider coarsened principal strata causal effects together with basic principal strata causal effects (Lee et al., 2010). However, as Robins (1986) noted, if one were to compare $E[Y(2) - Y(0) | G = LLL]$, $E[Y(1) - Y(0) | G = LLL]$ and $E[Y(2) - Y(1) | G \in \{LLL, DLL\}]$ simultaneously, it is possible that the last two comparisons are both positive while the first one is negative. This lack of transitivity limits the interpretability of causal effects defined within coarsened principal strata. In contrast, transitivity holds if limited to basic principal strata (e.g., *LLL*). Hence in this article, we are primarily interested in comparisons between potential outcomes in the same *basic* principal stratum.

On the other hand, as Robins et al. (2007) noted, the size of each basic principal stratum is likely to be very small and consequently, each comparison in (1) only applies to a small portion of the population. Hence for randomized trials with more than three treatment arms, we may have limited power to test treatment effects for each basic principal stratum. What is more, we run into the problem of multiple comparisons as there are multiple treatment arms and multiple basic principal strata.

Therefore, we first consider testing the global null hypothesis that the treatment is not effective in any of the basic principal strata (for which some treatment comparison is well-defined). This question is scientifically relevant. For example, in a HIV vaccine trial, testing the global null addresses whether there exists a mechanism through which the vaccine alters viral load in infected individuals (Shepherd et al., 2006). Secondly, clinicians may also be interested in whether the overall treatment effect is clinically meaningful so that the active treatment is promising in clinical practice. For this purpose, an overall treatment effect may be declared only if it is greater than the clinical margin of relevance specified by clinicians. Finally, besides an overall treatment effect, scientists and clinicians may also be interested in isolating the non-zero/non-trivial causal contrasts. In summary, the following questions are of interest with a multi-arm trial:

1. Is there evidence of the existence of *non-zero* average treatment effects for at least one basic principal stratum between at least two treatment arms?
2. Are there *clinically relevant* average treatment effects for at least one basic principal stratum between at least two treatment arms?
3. Can we find the specific principal strata and treatment arms that correspond to the overall nonzero/clinically relevant treatment effect, if such exists?

We address these questions in Section 3, 4 and 5 respectively. Existing causal analysis literature in multi-arm trials with non-identifiable causal estimands (Cheng and Small, 2006; Long et al., 2010; Lee et al., 2010) focuses on answering the third question. However, as we

explain later in Remark 3, one may be able to answer the first two questions even if there is not enough information to answer the third. Hence it is important to consider all three questions.

3 Testing treatment effects in a multi-arm trial

To find out if there is an overall non-zero treatment effect, it is desirable to consider the following testing problem:

$$\mathcal{H}_0: \Delta(z_1, z_2; g) = 0, \forall z_1, z_2, g \quad \text{vs} \quad \mathcal{H}_a: \exists z_1, z_2, g \text{ s.t. } \Delta(z_1, z_2; g) \neq 0, \quad (2)$$

where $\forall$ means “for all,” $\exists$ means “there exists” and *s.t.* means “such that.” The testing problem (2) is fundamentally different from (and more difficult than) a standard testing problem, in which one assumes if the observed data distribution was known, one would also know whether or not the hypothesis is true (Lehmann and Romano, 2006). The main difficulty here is that $\mathcal{H}_0$ is a statement about non-identifiable parameter vector μ_{m-1} . In other words, even if the population probabilities $P(S=1|Z=z)$ and $P(Y=1|S=1, Z=z)$ were known, we could only ascertain that μ_{m-1} resides in a region, and therefore may not know whether $\mathcal{H}_0$ is true or not.

Nevertheless, μ_{m-1} is “partially identifiable” in the sense that the observed data distribution can narrow down the range in which μ_{m-1} can possibly lie (Cheng and Small, 2006). For example, in a three-arm trial, the domain of μ_2 is $[0, 1]^6$. However, if the observed data distribution was known, the feasible region of μ_2 would be a subspace in $[0, 1]^6$ subject to the following constraints:

$$\begin{aligned} P(Y=1|Z=0, S=1) &= \mu_{LLL}^0, \\ P(Y=1|Z=1, S=1) &= p_{LLL}^1 \mu_{LLL}^1 + p_{DLL}^1 \mu_{DLL}^1, \\ P(Y=1|Z=2, S=1) &= p_{LLL}^2 \mu_{LLL}^2 + p_{DLL}^2 \mu_{DLL}^2 + p_{DDL}^2 \mu_{DDL}^2, \end{aligned} \quad (3)$$

where $p_g^z \equiv P(G=g|Z=z, S=1)$ is identifiable under Assumptions 2 and 3 (see Lemma 1 in the Supplementary Materials). Figure 1 provides a graphical representation of the functional relations described in (3).

For a general multi-arm trial, if the parameter space defined by $\mathcal{H}_0$ has no intersection with the feasible region of μ_{m-1} , one would know that $\mathcal{H}_0$ is not true. In general, we introduce the following notions for hypothesis testing with non-identifiable parameters.

Definition 1

We define a hypothesis relating to a parameter to be *compatible* with an observed data distribution if the parameter space defined by the hypothesis has a non-empty intersection with the feasible region of the parameter under the observed data distribution.

In particular, if a parameter is completely unidentifiable such that the observed data distribution imposes no constraints on the parameter, then all hypotheses relating to that parameter are compatible with the observed data distribution. On the other hand, if a parameter is identifiable so that its feasible region under the observed data distribution is always a single point set, then all compatible hypotheses are true.

In general, however, not all compatible hypotheses are true. Nevertheless, due to lack of identifiability, a true hypothesis may not be distinguishable from data with an untrue yet compatible hypothesis. This leads to the following notion of sharpness.

Definition 2

We define a test to be *sharp* for testing a null hypothesis if when the null is not compatible with the observed data distribution (and is hence untrue), the power of the test tends to 1 when the sample size goes to infinity.

Intuitively, similar to consistent tests, sharp tests are those that maximize power asymptotically. The difference is that as the sample size goes to infinity, with probability tending to 1, sharp tests reject any hypotheses that are incompatible with the observed data distribution, whereas consistent tests reject any hypotheses they are untrue. In small sample settings, however, the conclusions that one would draw from a sharp test are similar to those from a consistent test. If a hypothesis is rejected, one would conclude that it is untrue (at a certain significance level); if otherwise, no claims about the correctness of the hypothesis would be made. We also note that for a standard hypothesis testing problem as described in Lehmann and Romano (2006), sharp tests are the same as consistent tests. When the null hypothesis concerns nonidentifiable parameters, however, there are in general no consistent tests. Instead, sharpness plays the role of consistency in a standard hypothesis testing problem.

The notion of sharp tests is similar in spirit to the notion of sharp bounds, defined as the tightest possible bound given the observed data distribution (e.g., Imai, 2008). This notion has also been used implicitly in previous works. For example, Hudgens et al. (2003)'s test for SACE in a two-arm trial is sharp.

Below in Section 3.1, we develop a sharp test for problem (2) under the presumption that the observed data distribution is known. In other words, we assume the sample size is infinite such that there is no stochastic variation in the observed data. In Section 3.2 we incorporate sampling uncertainty to our proposed test using a Bayesian method.

3.1 A step-down procedure for testing the global null $\mathcal{H}_0$

To fix the idea, we first consider the problem for a three-arm trial, for which $\mathcal{H}_0$ holds if and only if

$$\mu_{LLL}^0 = \mu_{LLL}^1 = \mu_{LLL}^2 \quad (4)$$

and

$$\mu_{DLL}^1 = \mu_{DLL}^2. \quad (5)$$

We hence propose a two-step procedure. Firstly we test hypothesis (4). If (4) is compatible with the observed data distribution, we then test if (5) is compatible with the observed data distribution conditioning on (4).

Specifically, one can see from Figure 1 that μ_{LLL}^0 is identifiable from the observed data and suppose the feasible regions of μ_{LLL}^1 and μ_{LLL}^2 are B_{01} and B_{02} , respectively. If μ_{LLL}^0 is not contained in the intersection of B_{01} and B_{02} , then (4) and hence $\mathcal{H}_0$ is not compatible with the observed data distribution. If otherwise, that (4) is compatible with the observed data distribution, we then test hypothesis (5) conditioning on (4). Note under (4), μ_{LLL}^1 and μ_{LLL}^2 are identifiable. Consequently, μ_{LLL}^1 is identifiable. Suppose the feasible region of μ_{DLL}^2 under the constraint (4) is B_{12} . If μ_{DLL}^1 is not contained in B_{12} , we conclude that (5) is not compatible with the observed data distribution under the constraint (4) and hence reject $\mathcal{H}_0$. If otherwise, we conclude that $\mathcal{H}_0$ is compatible with the observed data distribution.

Algorithm 1 generalizes the procedure described above to general multi-arm trials. Theorem 1 states the asymptotic optimality of Algorithm 1. The proof is provided in the Supplementary Materials.

Algorithm 1

A step-down algorithm for testing the global null hypothesis $\mathcal{H}_0$

-
1. Set $k = 0$
 2. For $z = k, \dots, m$
 obtain the feasible region B_{kz} for $\mu_{D^k L^m+1-k}^z$ (see Theorem 2)
 3. If $\bigcap_{z=k, \dots, m} B_{kz} = \emptyset$
 reject $\mathcal{H}_0$; report k ; stop
 else
 set $\mu_{D^k L^m+1-k}^k = \dots = \mu_{D^k L^m+1-k}^m$ (6)
 4. If $k = m$
 fail to reject $\mathcal{H}_0$ and stop
 else
 set $k = k + 1$ and go to Step 2
-

Theorem 1—The test given by Algorithm 1 is sharp for testing $\mathcal{H}_0$. In other words, it is asymptotically optimal for testing $\mathcal{H}_0$ as it maximizes power given the observed data distribution.

To derive the feasible regions $(B_{kz}; k = 0, \dots, m, z = k, \dots, m)$ in Algorithm 1, we introduce notation building on Horowitz and Manski (1995). Let $Q_{\mathcal{G}}^z(\cdot)$ denote the distribution (function) of outcome Y among members of subgroup $\mathcal{G}$ who receive treatment z , and $\delta_x(\cdot)$ be a degenerate distribution function localized at x . As Y is binary, $Q_{\mathcal{G}}^z(\cdot)$ is a Bernoulli distribution with mean $m_{\mathcal{G}}(z)$: $Q_{\mathcal{G}}^z(\cdot) = (1 - m_{\mathcal{G}}(z)) \delta_0(\cdot) + m_{\mathcal{G}}(z) \delta_1(\cdot)$. To compress notation, we write $Q_{\mathcal{G}}^z(\cdot)$ as $Q_{\mathcal{G}}^z(\cdot)$. Also let $L_{\lambda}(Q)$ and $U_{\lambda}(Q)$ be functionals that map a distribution function Q to the corresponding distributions truncated at the lower λ quantile and upper λ quantile, respectively. Theorem 2 gives the formula for feasible region B_{Lz} . Intuitively, the bounds of B_{Lz} are obtained by assigning the smallest/largest ω_g^z portion of observed outcome values in distribution Q_g^z to principal stratum g . The proof is in the Supplementary Materials.

Theorem 2—Suppose that the observed data distribution is known and (6) holds for all $k < l$. Let $g = D^l L^{m+1-1}$ and $\bar{g} = \bigcup_{\bar{g} \in \Omega_z \setminus \Omega_{l-1}} \bar{g}$ be the coarsened principal stratum whose members would survive if assigned to treatment z but would die if assigned to treatment $l-1$. The feasible region of μ_g^z is

$$B_{Lz} = \left[\int y dL_{\omega_g^z}(Q_g^z), \int y dU_{\omega_g^z}(Q_g^z) \right], \quad (7)$$

where $\omega_g^z \equiv P[G=g|G \in \Omega_z \setminus \Omega_{l-1}] = p_g^z / \left(\sum_{\bar{g} \in \Omega_z \setminus \Omega_{l-1}} p_{\bar{g}}^z \right)$ and Q_g^z is a Bernoulli distribution with mean

$$m_g(z) = \left(m(z) - \sum_{\bar{g} \in \Omega_{l-1}} p_{\bar{g}}^z \mu_{\bar{g}}^z \right) / \left(1 - \sum_{\bar{g} \in \Omega_{l-1}} p_{\bar{g}}^z \right),$$

in which $m(z) \equiv P[Y=1|Z=z, S=1]$.

Remark 1—Algorithm 1 is a “step-down” procedure in the sense that the hypothesis $\mathcal{H}_0$ is decomposed into a series of hypotheses where the first hypothesis concerns the first stratum L^{m+1} , the second hypothesis concerns the second stratum DL^m conditioning on the first hypothesis, and so on.

3.2 Bayesian procedures

We have so far developed a sharp test for problem (2). In practice, however, sampling uncertainty must be taken into account when making statistical inference. Here we introduce a Bayesian procedure to estimate the posterior probability that $\mathcal{H}_0$ is not compatible with the observed data distribution. The Bayesian method produces multiple samples of the posterior distribution, thereby reflecting randomness in observed data.

Let $p(s, y|z) = P(S=s, Y=y|Z=z)$ and $p(\cdot, \cdot|z) = (p(1, 1|z), p(1, 0|z), p(0, \cdot|z))$, where $\uparrow$ indicates that Y is undefined when $S=0$. Define $\mathbf{p} = (p(\cdot, \cdot|0), \dots, p(\cdot, \cdot|m))$. Under

independent Dirichlet priors over the *observed distributions* $p(\cdot, \cdot | z)$, $z = 0, \dots, m$, it is easy to sample from the posterior distribution via conjugacy. We propose to use Algorithm 2 to calculate the posterior probability that $\mathcal{H}_0$ is not compatible with the observed data distribution.

Remark 2—The step-down procedure in Algorithm 1 has a similar structure to the sequential tests for nested hypotheses discussed by Rosenbaum (2008). His procedure has attractive Frequentist properties since it controls the type I error rate without resorting to multiplicity adjustment. However, with his methods one proceeds to the next step if the current hypothesis is rejected whereas in our proposal, one proceeds if the current hypothesis is *not* rejected. Moreover, in his context, the parameters of interest are identifiable. Hence his results are not directly applicable to our case.

Algorithm 2

A Bayesian procedure for testing $\mathcal{H}_0$

-
1. Place an independent Dirichlet prior $Dir(a_{3z+1}; a_{3z+2}; a_{3z+3})$ on $p(\cdot, \cdot | z)$, $z = 0, \dots, m$.
 2. Simulate samples $p^{(1)}, \dots, p^{(M)}$ from the posterior distributions, which are independent Dirichlet distributions

$$Dir(a_{3z+1} + n_{3z+1}, a_{3z+2} + n_{3z+2}, a_{3z+3} + n_{3z+3}), z = 0, \dots, m,$$

where

$$n_{3z+1} = \sum_{i=1}^N I(S_i = 1, Y_i = 1, Z_i = z), n_{3z+2} = \sum_{i=1}^N I(S_i = 1, Y_i = 0, Z_i = z), n_{3z+3} = \sum_{i=1}^N I(S_i = 0, Z_i = z)$$

3. Run Algorithm 1 with each of the posterior samples satisfying the following inequalities:

$$P(S = 1 | Z = m) \geq \dots \geq P(S = 1 | Z = 1) \geq P(S = 1 | Z = 0) \quad (8)$$

Note (8) characterizes the set of observed data distributions arising from the potential outcome model defined by Assumptions 1 - 3.

4. Report the proportion of posterior samples with which H_0 is rejected.
-

4 Testing clinically relevant treatment effects in a multi-arm trial

If a non-zero treatment effect is found using Algorithm 2, a natural question arises that whether the treatment effect is clinically meaningful. Suppose the margin of clinical relevance is Δ_0 such that a treatment effect smaller than this would not matter in practice, and also suppose that the treatment effect is clinically meaningful only if a higher treatment level corresponds to a higher mean potential outcome. It is desirable to consider the following testing problem:

$$\mathcal{H}_{0,c}: \Delta(z_1, z_2; g) \leq \Delta_0, \forall g, z_1 \geq z_2 \quad \text{vs} \quad \mathcal{H}_{a,c}: \exists g, z_1 \geq z_2 \text{ s.t. } \Delta(z_1, z_2; g) > \Delta_0, \quad (9)$$

where the letter “ c ” in $\mathcal{H}_{0,c}$ is short for “clinical relevance.” Similar to (2), (9) is a testing problem on non-identifiable parameters. However, as the null parameter space is a non-degenerate region in the domain of μ_{m-1} , the step-down procedure developed in Section 3 is

not applicable. Instead, we define Δ_{max} to be the largest $\Delta(z_1, z_2; g)$ that appears in $\mathcal{H}_{0,c} : \Delta_{max} = \max_{g, z_1 \geq z_2} \Delta(z_1, z_2; g)$. (9) can then be rewritten in an *equivalent* form using $\Delta_{max} : \mathcal{H}_{0,c} : \Delta_{max} \leq \Delta_0$ vs $\mathcal{H}_{\alpha,c} : \Delta_{max} > \Delta_0$. The following lemma says the testing problem (9) can be translated into the identification problem on Δ_{max} .

Lemma 3

Suppose the sharp (large sample) lower bound for Δ_{max} is $\Delta_{max,slb}$. A sharp test would reject $\mathcal{H}_{0,c}$ if and only if $\Delta_{max,slb} > \Delta_0$.

As Δ_{max} is a function of μ_{m-1} , in general, identifying $\Delta_{max,slb}$ involves minimizing Δ_{max} subject to the constraints on μ_{m-1} imposed by the observed data distribution. Theorem 4 below says that the feasible region of μ_{m-1} is a convex polytope, defined as an intersection of finitely many half spaces. Consequently, this optimization problem can be translated into a linear programming problem and efficiently solved with off-the-shelf software. See Algorithm 1 in the Supplementary Materials for more details.

Theorem 4

Given the observed data distribution, the feasible region of μ_{m-1} is a subspace in $[0,1]^{dim(\mu_{m-1})}$ subject to the following constraints:

$$\sum_{g \in \Omega_z} p_g^z \mu_g^z = m(z), z=0, \dots, m-1;$$

$$\max \left(0, m(z) - p_{D^m L}^z \right) \leq \sum_{g \in \Omega_{m-1}} p_g^z \mu_g^z \leq \min \left(1 - p_{D^m L}^z, m(z) \right), z=m,$$

where p_g^z is identifiable from data under Assumptions 2 and 3 (see Lemma 1 in the Supplementary Materials). In particular, the feasible region of μ_{m-1} is a convex polytope.

To incorporate statistical uncertainty, one can use Bayesian analysis methods to derive a credible interval for $\Delta_{max,slb}$. Specifically, one runs Steps 1-4 in Algorithm 2 to get multiple posterior samples that satisfy the constraint (8), and then produces a percentile based credible interval for $\Delta_{max,slb}$ based on the posterior samples. One may also estimate the posterior probability of rejecting $\mathcal{H}_{0,c}$ for any given positive value Δ_0 with these posterior sample draws.

5 Marginal credible intervals for a given contrast

If a clinically non-trivial treatment effect is found, then it is desirable to identify the principal strata and treatment arms that correspond to this treatment effect. In this case, the marginal feasible regions and associated credible intervals for $\Delta(z_1, z_2; g)$ are of interest.

If the observed data distribution was known, then the feasible region for $\Delta(z_1, z_2; g)$ can be obtained from the feasible regions for $\mu_g^{z_1}$ and $\mu_g^{z_2}$. Specifically, we have the following theorem.

Theorem 5

Suppose the observed data distribution is known, and $B_{\mathcal{M}(g), z_1}$ and $B_{\mathcal{M}(g), z_2}$ are feasible regions for $\mu_g^{z_1}$ and $\mu_g^{z_2}$, respectively. Then we have the following results.

1. For $z = z_1, z_2$, $B_{\mathcal{M}(g), z} = \left[\int y dL_{p_g^z}(Q^z), \int y dU_{p_g^z}(Q^z) \right]$.
2. The feasible region of $(z_1, z_2; g)$ is $\left[\int y dL_{p_g^{z_1}}(Q^{z_1}) - \int y dU_{p_g^{z_2}}(Q^{z_2}), \int y dU_{p_g^{z_1}}(Q^{z_1}) - \int y dL_{p_g^{z_2}}(Q^{z_2}) \right]$.

In practice, credible intervals for $(z_1, z_2; g)$ can be constructed from posterior sample draws $p^{(1)}, \dots, p^{(M)}$. These posterior draws may also be used to estimate the posterior probability of rejecting the null hypothesis $\mathcal{H}_{0,m}: \Delta(z_1, z_2; g) \leq \Delta_0$, where the letter “ m ” in $\mathcal{H}_{0,m}$ is short for “marginal.”

Remark 3

We remark that even if the observed data provide evidence for existence of non-zero/non-trivial treatment effects, it is possible that they do not contain information on the specific principal strata and treatment arms that correspond to these treatment effects. Moreover, unlike the case in multi-arm trials without truncation by death, this can happen even with an infinite sample size.

We illustrate our point with the following numerical example. Consider a three-arm trial such that $\pi_{LLL} = \pi_{DLL} = \pi_{DDL} = 0.3$, $\pi_{DDD} = 0.1$, $m(0) = 0.3$, $m(1) = 0$, $m(2) = 0.5$, where $\pi_g \equiv P(G = g)$. In this case, $\mu_{LLL}^0 = 0.3$ and $\mu_{LLL}^1 = \mu_{DLL}^1 = 0$. It follows that

$\Delta_{max} = \max(0, \mu_{LLL}^2 - \mu_{LLL}^1, \mu_{DLL}^2 - \mu_{DLL}^1)$. We assume that the sample size is infinite so that we know the observed data distribution. Figure 2 shows the joint feasible region of $(\mu_{LLL}^2 - \mu_{LLL}^1, \mu_{DLL}^2 - \mu_{DLL}^1)$ (the green shaded area). Suppose that the margin of clinical relevance Δ_0 is 0.1, then the acceptance region for null hypothesis $\mathcal{H}_{0,c}$ is the lower left area of the blue contour line. As there is no intersection between the feasible region of $(\mu_{LLL}^2 - \mu_{LLL}^1, \mu_{DLL}^2 - \mu_{DLL}^1)$ and the acceptance region for $\mathcal{H}_{0,c}$, one may conclude that $\mathcal{H}_{0,c}$ should be rejected. Alternatively, one can see from the contour lines of Δ_{max} that the sharp lower bound of Δ_{max} is 0.25. As Δ_0 is smaller than $\Delta_{max,slb}$, one also rejects $\mathcal{H}_{0,c}$.

However, by projecting the joint feasible region of $(\mu_{LLL}^2 - \mu_{LLL}^1, \mu_{DLL}^2 - \mu_{DLL}^1)$ onto individual axes, one concludes that the marginal feasible regions for $\mu_{LLL}^2 - \mu_{LLL}^1$ and $\mu_{DLL}^2 - \mu_{DLL}^1$ are both $[0, 1]$. As both of the marginal feasible regions contain values that are smaller than Δ_0 , the data contain no information on the specific contrast that corresponds to the overall treatment effect.

6 Data Illustrations

6.1 Application to the HVTN 503 study

The HVTN 503 HIV vaccine study was a randomized, double-blinded, placebo-controlled Phase IIb test-of-concept clinical trial to investigate the efficacy and safety of an experimental HIV vaccine. The same vaccine was also evaluated in a different population in an earlier HVTN 502/Step trial. Starting January, 2007, the HVTN 503 study enrolled 800 HIV negative subjects and randomized them to receive three doses of either the study vaccine or a placebo. The ratio of vaccine to placebo assignment was 1:1. Enrollment and vaccinations were halted in September 2007, but follow-up continued, after the HVTN 502/Step trial met its prespecified non-efficacy criteria. Details of this study can be found in Gray et al. (2011, 2014).

In our analysis, we compared CD4 counts among participants within the same principal stratum defined by their full potential infection statuses. Due to the early stopping of vaccinations of the trial, a majority of participants in the HVTN 503 trial were not fully immunized. When enrollment was stopped, 400 participants in the HVTN 503 trial were assigned to the experimental vaccine group. Of them, 112 received one injection, 259 received two injections, and only 29 received all three injections. Hence we considered the dosage of experimental vaccine as the treatment arm Z , where $Z = 0$ for all subjects in the control group. As the trial was stopped administratively, and the time a participant entered this trial was unlikely to affect the potential outcomes of interest (CD4 count), it is reasonable to assume that the treatment arms were randomized. Furthermore, since there were only 3.6% of participants who received all three experimental vaccines, we code $Z = 2$ for all participants who receive two or more experimental vaccine injections.

A total of 100 subjects were infected during this trial. We defined each subject's "median CD4 count" (the outcome of interest) as their median CD4 count measured between their confirmatory HIV testing visit and the end of follow-up or start of antiretroviral treatments. We also dichotomized CD4 count at 350 cells/mm³ and 200 cells/mm³ as they have been used in previous United States Department of Health and Human Services (DHHS) guidelines for initiating antiretroviral treatment. Note that the outcome measure is only measured for infected subjects. As 87.5% of the study subjects were uninfected, an intent-to-treat analysis with imputation for missing CD4 count values is likely to have very low power for detecting any treatment effects (Gilbert et al., 2003). Hence SACEs are of interest for analyzing this trial.

Table 2 in the Supplementary Materials summarizes the observed data for the study participants. There were 7 infected participants who had no CD4 count measurements after their confirmatory HIV testing visit. We made the missing completely at random (MCAR) assumption and left them out of our analysis below. In treatment arm 0, 1, 2, the mean number of CD4 counts available were 5.69, 5.94 and 5.57, respectively; the mean length of time from the confirmatory HIV testing visit to the first CD4 count measure were 26 days, 25 days and 32 days, respectively, and the mean time spacing between CD4 count measurements were 127 days, 146 days and 134 days, respectively.

Presumably there was little interaction among HVTN 503 subjects so that the SUTVA was plausible. Subsequent analyses of the HVTN 502/503 data suggested that although not possible to directly cause HIV infections itself, the investigational vaccine may increase susceptibility to HIV infection for recipients (Gray et al., 2011, 2014). Given the negative results on the primary efficacy endpoints, members of the HVTN 503 Protocol Team whom we consulted agreed that it is reasonable to make the reverse monotonicity assumption such that experimental vaccine did not help prevent HIV infection for any participant in the study population. The empirical infection rates in the $Z = 0, 1, 2$ arms were 9.25%, 16.07% and 15.63%, respectively. Thus, the reverse monotonicity assumption seemed acceptable, and we proceeded with our analysis under the reverse monotonicity assumption.

Table 1 summarizes the analysis results. The simultaneous testing method estimates the posterior probability of existence of an overall non-zero treatment effect, while the marginal testing method estimates the posterior probability that an overall non-zero treatment effect can be claimed along with the specific treatment arms and principal strata that correspond to this treatment effect. These posterior probabilities were high, suggesting evidence of a non-zero treatment effect on median CD4 falling below 350 or 200 cells/mm³. The 95% credible intervals for lower bound on $\max$ provide information on the magnitude of vaccine effects. For example, results in Tables 1 show that there exists at least one basic principal stratum and treatment comparison for which the vaccine reduces the probability of median CD4 count ≥ 200 cells/mm³ by 0.026, but we were not able to ascertain the specific basic principal stratum and treatment comparison that corresponds to this effect. The reason for this is two fold. Firstly, due to non-identifiability of the SACEs, if the effect size is too small, one may fail to identify the specific causal contrast that corresponds to a clinically relevant treatment effect even with an infinite sample size. Secondly, our proposed methods may deliver more conclusive results if the sample size is large enough. For example, if the sample size was 3000 (which was the estimated sample size in the HVTN 503 trial protocol) and the observed frequencies $P(S = 1|Z = z)$ and $P(Y = 1|Z = z, S = 1)$ had remained the same, then the 95% credible interval for the contrast $\mu_{LLL}^2 - \mu_{LLL}^0$ would have been [0.057, 0.186], which would imply that compared to the placebo, receiving two or more injections of the experimental vaccine is clinically effective for reducing the possibility of very low CD4 counts (< 200 cells/mm³) among subjects who would get infected regardless of which treatment arm they were assigned to.

We conclude this part with several caveats. First, the median CD4 count is a non-traditional endpoint for HIV vaccine efficacy trials, and it may not be completely comparable between treatment groups due to differences in the number and timing of CD4 measurements. Second, we have dichotomized CD4 count in our analysis, which results in loss of information. Third, we have made the MCAR assumption for the missing values in CD4 count measures, which is hard to verify for this data set. Fourth, as pointed out by some authors (e.g. Pearl, 2011), under the principal stratification framework we have taken here, the vaccine effect estimates are only relevant for the subgroup of subjects who would get infected under at least one dosage level, which only constitutes a small fraction of the population. Finally, a reduction of 0.026 in the probability of median CD4 counts ≥ 200

cells/mm³ may not be considered clinically important given the earlier finding that the vaccine increased HIV acquisition in the study population.

6.2 Application to survey incentive trials

Faced with declining voluntary participation rates, there is now a consensus that incentives are effective for motivating response to surveys (Singer and Kulka, 2002; Singer and Ye, 2013). There is, however, controversy on how incentives affect the quality of data collected. Social exchange theory suggests that by establishing an explicit exchange relationship, incentives not only encourage participation in surveys, but also encourage respondents to provide more accurate and complete information (Davern et al., 2003). However, current experimental studies have mixed findings on this hypothesis (Singer and Kulka, 2002; Singer and Ye, 2013).

These experimental studies directly compare response quality in different incentive groups without accounting for the problem of truncation by response. Here the treatments Z are the levels of incentive, the intermediate outcomes S are the responses to the surveys, and the final outcomes Y are measures of survey quality. Although some researchers realize that people persuaded to participate through the use of incentives will have less internal motivation for filling out the survey thoroughly (e.g. Davern et al., 2003), few, if any, of them separate this group of people in their analyses from those who would participate in the survey regardless of incentive levels, rendering their results subject to selection bias. Furthermore, arguably the response quality is *undefined* for survey non-respondents. Thus as argued by Rubin (2006) and others, the naive comparison is not causal as it compares different groups of people at baseline. Instead, for two-arm trials, the SACE is of interest as the subgroup whose members would respond regardless of incentive level is the only group for which both of the potential outcomes are well-defined. Similarly for multi-arm trials. Moreover, it is very common that such randomized experiments have multiple incentive groups (Singer and Kulka, 2002; Singer and Ye, 2013). Hence the methodology introduced in this paper, and more generally, identification and estimation methods for SACEs in multi-arm trials are especially relevant.

For example, Curtin et al. (2007) used data from the Survey of Consumer Attitudes (SCA) conducted by the University of Michigan Survey Research Center to investigate whether efforts to increase the response rate jeopardize response quality. Their analysis was based on a random digit dial telephone survey conducted between November 2003 and February 2004. In each of the four months, eligible samples were randomly assigned to one of three experimental conditions: advance letter without an incentive, advance letter plus \$5 incentive and advance letter plus \$10 incentive. The same follow-up procedures, including promised refusal conversion payments are used in all three groups. The measure for response quality in such studies are inevitably subjective; they can be binary (e.g., “mostly compete” vs “partially complete,” or whether a particularly important question is answered) or continuous (e.g. percent of missing items). As we don’t have access to this data set, below we only discuss the validity of our assumptions.

The SUTVA is reasonable as these are random digit dial samples from the coterminous United States. The monotonicity assumption is also plausible. As argued by survey sampling

experts, incentives will motivate response as they compensate for the relative absence of factors that might otherwise stimulate cooperation (Singer and Kulka, 2002), so that individuals who would respond with a lower incentive would also respond if offered a higher incentive. Empirical evidence in this study also supports this assumption: the response rate for the three experimental groups were 51.7%, 63.8% and 67.7% (Curtin et al., 2007).

7 Discussion

In randomized trials with truncation by death, the average causal effects in basic principal strata are often of interest as they provide causally meaningful and interpretable summaries of the treatment effects. However, for trials with multiple treatment arms, there are usually many such causal contrasts that are of interest to investigators. In this article, we consider testing and estimation problems on the basic principal stratum causal effects. Specifically, we propose three scientific questions to understand the overall treatment effect and individual principal stratum causal effects. We then develop novel inference procedures to answer these questions, and show that the proposed procedures have desirable asymptotic properties.

Compared to analyzing a multi-arm trial in a standard setting, the main difficulty introduced by truncation by death is that the causal estimands are not identifiable. In this case, we show that compared to marginal methods, the (ANOVA type) simultaneous inference methods provide more power for testing the overall treatment effect, and the advantage remains even with an infinite sample size. These results demonstrate the importance of addressing both joint and marginal hypotheses in a causal analysis of multi-arm trials with truncation by death. This idea may be applied to analyze multi-arm trials in other settings in which the causal estimands are not identifiable. For example, in multi-arm trials with non-compliance, existing methods consider the causal contrasts separately (Cheng and Small, 2006; Long et al., 2010). Although results obtained with such methods are valid, they are often not informative, especially in the case where there are more than three treatment arms (Long et al., 2010). In this case, a simultaneous inference method may yield a greater posterior probability of claiming an overall treatment effect and the posterior credible intervals are less likely to contain zero.

In analyzing a multi-arm trial with truncation by death, researchers may dichotomize the treatment variable to simplify an analysis, especially in settings where the multi-arm trials consist of a placebo arm and several dosage groups for an active treatment. One such example is the HVTN 503 study, where the treatment groups 1 and 2 can be considered as different versions of the experimental vaccine. However, as noted by Hernán and VanderWeele (2011), results from these analyses may not be generalizable to other population as the causal effect of a compound treatment depends on the distribution of treatment versions in the target population. Moreover, due to the non-identifiability of SACEs, one may fail to find an overall treatment effect that could have been found by applying the proposed simultaneous inference procedure. For example, for the HVTN 503 study, if one were to collapse the active treatment groups into a single compound treatment, then the 95% credible intervals for the SACE corresponding to this compound treatment

would be [0.000, 0.253], with which one could not claim any clinically relevant treatment effect.

To account for sampling uncertainty in the observed data distribution, we use Bayesian analysis methods to obtain posterior samples of identifiable quantities p . An alternative Bayesian procedure to our method involves posterior sampling on the mean potential outcomes μ_{m-1} . This alternative approach would directly yield the posterior rejection rate of $\mathcal{H}_0$ and credible intervals for $\max, slb$ without resorting to techniques we have introduced. However, as μ_{m-1} is not identifiable from the observed data, it turns out that the posterior estimates of $\mu_{\max}$ are extremely sensitive to the prior specification on μ_{m-1} . We refer interested readers to Richardson et al. (2011) for a nice discussion of this issue.

The problem we consider here is similar to an instrumental variable analysis in that both problems can be analyzed under the principal stratification framework. When the exposure variable in a instrumental variable analysis is binary, the exclusion restriction assumption is closely related to the null hypothesis in the truncation by death problem, namely the causal effect in the always-survivor group is zero. Hence the approach we develop here may be used to partially test the exclusion restriction assumption of an instrumental variable model.

There are several possible extensions to our framework. For example, we have restricted our attention to binary outcomes in this article. We are currently exploring extensions to deal with continuous and categorical outcomes. In addition, covariate information may be employed to sharpen bounds on SACEs. Another possible extension is to introduce sensitivity parameters for better understanding of the causal effects of interest. The tests and bounds we have developed here correspond to extreme results of corresponding sensitivity analyses.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

Research reported in this paper was supported by the National Institute of Allergy and Infectious Diseases (NIAID) of the National Institutes of Health (NIH) under Award Number UM1AI068635. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH or the Department of Veterans Affairs. The authors thank the NIAID-funded HIV Vaccine Trials Network for providing the dataset from the HVTN 503 trial. Furthermore, the authors thank the participants, investigators, and sponsors of the HVTN 503 trial. The authors also thank Peter Gilbert and members of the HVTN Ancillary Study Committee for valuable comments. Richardson was supported by National Institutes of Health Grant R01 AI032475. Zhou was supported in part by a US Department of Veterans Affairs, Veterans Affairs Health Administration, Research Career Scientist award (RCS 05-196).

References

- Cheng J, Small DS. Bounds on causal effects in three-arm trials with non-compliance. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2006; 68(5):815–836.
- Chiba Y, VanderWeele TJ. A simple method for principal strata effects when the outcome has been truncated due to death. *American journal of epidemiology*. 2011; 173(7):745–751. [PubMed: 21354986]

- Curtin R, Singer E, Presser S. Incentives in random digit dial telephone surveys: A replication and extension. *Journal of Official Statistics*. 2007; 23(1):91–105.
- Davern M, Rockwood TH, Sherrod R, Campbell S. Prepaid monetary incentives and data quality in face-to-face interviews: Data from the 1996 survey of income and program participation incentive experiment. *Public Opinion Quarterly*. 2003; 67(1):139–147.
- Ding P, Geng Z, Yan W, Zhou X-H. Identifiability and estimation of causal effects by principal stratification with outcomes truncated by death. *Journal of the American Statistical Association*. 2011; 106(496):1578–1591.
- Egleston BL, Scharfstein DO, Freeman EE, West SK. Causal inference for non-mortality outcomes in the presence of death. *Biostatistics*. 2007; 8(3):526–545. [PubMed: 16980696]
- Frangakis CE, Rubin DB. Principal stratification in causal inference. *Biometrics*. 2002; 58(1):21–29. [PubMed: 11890317]
- Gilbert PB, Bosch RJ, Hudgens MG. Sensitivity analysis for the assessment of causal vaccine effects on viral load in HIV vaccine trials. *Biometrics*. 2003; 59(3):531–541. [PubMed: 14601754]
- Gray GE, Allen M, Moodie Z, Churchyard G, Bekker L-G, Nchabeleng M, Mlisana K, Metch B, de Bruyn G, Latka MH, et al. Safety and efficacy of the HVTN 503/Phambili study of a clade-B-based HIV-1 vaccine in South Africa: a double-blind, randomised, placebo-controlled test-of-concept phase 2b study. *The Lancet infectious diseases*. 2011; 11(7):507–515. [PubMed: 21570355]
- Gray GE, Moodie Z, Metch B, Gilbert PB, Bekker L-G, Churchyard G, Nchabeleng M, Mlisana K, Laher F, Roux S, et al. Recombinant adenovirus type 5 HIV gag/pol/nef vaccine in South Africa: unblinded, long-term follow-up of the phase 2b HVTN 503/Phambili study. *The Lancet infectious diseases*. 2014; 14(5):388–396. [PubMed: 24560541]
- Hernán MA, VanderWeele TJ. Compound treatments and transportability of causal inference. *Epidemiology*. 2011; 22(3):368–377. [PubMed: 21399502]
- Horowitz J, Manski CF. Identification and robustness with contaminated and corrupted data. *Econometrica*. 1995; 63(2):281–302.
- Hudgens MG, Hoering A, Self SG. On the analysis of viral load endpoints in HIV vaccine trials. *Statistics in Medicine*. 2003; 22(14):2281–2298. [PubMed: 12854093]
- Imai K. Sharp bounds on the causal effects in randomized experiments with truncation-by-death. *Statistics & probability letters*. 2008; 78(2):144–149.
- Kalbfleisch, J., Prentice, R. The statistical analysis of failure time data. Wiley series in probability and mathematical statistics. 1980.
- Lee K, Daniels MJ, Sargent DJ. Causal effects of treatments for informative missing data due to progression/death. *Journal of the American Statistical Association*. 2010; 105(491):912–929. [PubMed: 21318119]
- Lehmann, EL., Romano, JP. Testing statistical hypotheses. Springer Science & Business Media; 2006.
- Long Q, Little RJ, Lin X. Estimating causal effects in trials involving multitreatment arms subject to non-compliance: a Bayesian framework. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*. 2010; 59(3):513–531.
- Pearl J. Principal stratification—a goal or a tool? *The International Journal of Biostatistics*. 2011; 7(1): 1–15.
- Richardson TS, Evans RJ, Robins JM. Transparent parameterizations of models for potential outcomes. *Bayesian Statistics*. 2011; IX:569–610.
- Robins J. A new approach to causal inference in mortality studies with a sustained exposure period – application to control of the healthy worker survivor effect. *Mathematical Modelling*. 1986; 7(9): 1393–1512.
- Robins J, Rotnitzky A, Vansteelandt S, Frangakis CE, Rubin DB, An M-W, MacKenzie E. Discussion of Principal stratification designs to estimate input data missing due to death. *Biometrics*. 2007; 63(3):650–653. [PubMed: 17824996]
- Rosenbaum PR. Testing hypotheses in order. *Biometrika*. 2008; 95(1):248–252.
- Rubin DB. Comment. *Journal of the American Statistical Association*. 1980; 75(371):591–593.
- Rubin DB. Causal inference without counterfactuals: comment. *Journal of the American Statistical Association*. 2000; 95(450):435–438.

- Rubin DB. Causal inference through potential outcomes and principal stratification: Application to studies with “censoring” due to death. *Statistical Science*. 2006; 21(3):299–309.
- Schulz KF, Grimes DA. Multiplicity in randomised trials I: endpoints and treatments. *The Lancet*. 2005; 365(9470):1591–1595.
- Shepherd BE, Gilbert PB, Jemai Y, Rotnitzky A. Sensitivity analyses comparing outcomes only existing in a subset selected post-randomization, conditional on covariates, with application to HIV vaccine trials. *Biometrics*. 2006; 62(2):332–342. [PubMed: 16918897]
- Singer, E., Kulka, RA. Paying respondents for survey participation. In: Ploeg, MV.Moffitt, RA., Citro, CF., editors. *Studies of welfare populations: Data collection and research issues*. National Academy Press; Washington, DC: 2002. p. 105-128.
- Singer E, Ye C. The use and effects of incentives in surveys. *The Annals of the American Academy of Political and Social Science*. 2013; 645(1):112–141.
- Zhang JL, Rubin DB. Estimation of causal effects via principal stratification when some outcomes are truncated by “death”. *Journal of Educational and Behavioral Statistics*. 2003; 28(4):353–368.

Figure 1.

A graph representing the functional dependencies in the causal analysis of a three-arm randomized trial with truncation by death. Rectangular nodes represent observed variables; oval nodes represent unknown parameters, with different shadings corresponding to different principal strata. Under the monotonicity assumption, p_g^z can be identified from observed quantities $P(S=1|Z=z)$.

Figure 2.

Feasible region of $(\mu_{LLL}^2 - \mu_{LLL}^1, \mu_{DLL}^2 - \mu_{DLL}^1)$ (green shaded area). The colored lines are contour lines of $\max$. The sharp lower bound of $\max$ is obtained at the red point.

Table 1

Posterior probabilities of finding a non-zero overall treatment and posterior credible intervals for lower bounds on τ_{max} (the maximal treatment effect over all principal strata and treatment comparisons) for the HVTN 503 trial

Methods	Posterior probability of a non-zero treatment effect	95% credible interval for lower bound on τ_{max}
Outcome: median CD4 > 350		
Simultaneous	0.882	[0.000, 0.346]
Marginal	0.651	[0.000, 0.341]
Outcome: median CD4 > 200		
Simultaneous	0.996	[0.026, 0.260]
Marginal	0.973	$[6 \times 10^{-4}, 0.245]$