

Identification of Vaccine Effects When Exposure Status Is Unknown

 Mats J. Stensrud^a and Louisa Smith^{b,c}

Results from randomized controlled trials (RCTs) help determine vaccination strategies and related public health policies. However, defining and identifying estimands that can guide policies in infectious disease settings is difficult, even in an RCT. The effects of vaccination critically depend on characteristics of the population of interest, such as the prevalence of infection, the number of vaccinated, and social behaviors. To mitigate the dependence on such characteristics, estimands, and study designs, that require conditioning or intervening on exposure to the infectious agent have been advocated. But a fundamental problem for both RCTs and observational studies is that exposure status is often unavailable or difficult to measure, which has made it impossible to apply existing methodology to study vaccine effects that account for exposure status. In this study, we present new results on this type of vaccine effects. Under plausible conditions, we show that point identification of certain relative effects is possible even when the exposure status is unknown. Furthermore, we derive sharp bounds on the corresponding absolute effects. We apply these results to estimate the effects of the ChAdOx1 nCoV-19 vaccine on SARS-CoV-2 disease (COVID-19) conditional on postvaccine exposure to the virus, using data from a large RCT.

Keywords: Estimands; Randomized Trials; Sharp bounds; Vaccine effects

(*Epidemiology* 2023;34: 216–224)


Submitted April 22, 2022; accepted November 28, 2022

From the ^aDepartment of Mathematics, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland; ^bDepartment of Health Sciences, Bouvé College of Health Sciences, Northeastern University, Boston, MA; ^cRoux Institute, Northeastern University, Portland ME.

M.J.S. is funded by the Swiss National Science Foundation, Grant 200021_207436.

The author reports no conflicts of interest.

Availability of data and code for replication: Code to replicate Figure 3 is available in the eAppendix, <http://links.lww.com/EDE/B994>.

 Supplemental digital content is available through direct URL citations in the HTML and PDF versions of this article (www.epidem.com).

Correspondence: Mats J. Stensrud, Department of Mathematics, Ecole Polytechnique Fédérale de Lausanne, EPFL SB MATH BIostat, MA B2 477 (Bâtiment MA), Station 8, CH-1015 Lausanne, Switzerland. E-mail: mats.stensrud@epfl.ch.

Copyright © 2023 The Author(s). Published by Wolters Kluwer Health, Inc. This is an open-access article distributed under the terms of the Creative Commons Attribution-Non Commercial-No Derivatives License 4.0 (CCBY-NC-ND), where it is permissible to download and share the work provided it is properly cited. The work cannot be changed in any way or used commercially without permission from the journal.

ISSN: 1044-3983/23/342-216-224

DOI: 10.1097/EDE.0000000000001573

Vaccines are one of the most important inventions in modern medicine.¹ Justification for real-life vaccination strategies relies heavily on results from large-scale vaccine randomized controlled trials (RCTs). However, the nature of communicable diseases means that defining and evaluating vaccine effects requires consideration of population characteristics such as the prevalence of current and prior infection, mixing patterns, and concurrent public health measures.

Policy-relevant estimands for vaccine trials have been discussed extensively (see Halloran et al.² for an overview), in particular in the context of the SARS-CoV-2 disease (COVID-19) pandemic.^{3–9} However, as of yet, methods to study vaccine effects conditional on, or under interventions on, exposure to the infectious agent are rarely used. Here and henceforth, we use exposure to mean exposure to the disease agent, such as a virus, which is distinct from the treatment, such as a vaccine. A key problem is that exposure status is often difficult, or even impossible, to measure in practice.^{2,10} For example, Halloran and Struchiner¹¹ write that measuring susceptibility to infection “might not be easy in practice and might indeed require considerable assumptions regarding who is infectious and when, how infectious the persons are, and who is exposing whom.” Challenge trials, in which participants are intentionally exposed, are one option for controlling exposure status but involve serious ethical issues.^{12–14}

This article specifically targets effects that account for exposure status, even when it is unmeasured. We provide results on the interpretation and identification of causal effects of vaccines from RCTs and observational studies. The results include identification conditions and formulas for the causal effect of a vaccine on clinical outcomes, conditional on an unmeasured exposure to the infectious agent. Specifically, we show that, under a plausible no effect on exposure assumption, the relative effect—though not the absolute effect—of the vaccine can be point-identified in an RCT. Furthermore, under the same assumption, we derive sharp bounds for the absolute effect. We clarify how these effects are related to existing estimands and notions of biological effects, and we give identification results on per-exposure effects,^{10,15} a type of controlled direct effect, even when the exposure is unmeasured, as is often the case in practice.

The article is organized as follows. Section Data Structure presents the data structure and the notation. Section Causal Parameters provides definitions and interpretation of

causal estimands. Section Identification contains results on the identification of causal estimands, including point identification results for the relative causal effect conditional on exposure, and partial identification results for the absolute causal effect conditional on exposure. Section External Data and Sensitivity Analysis presents results for point identification of absolute causal effects conditional on exposure when external data on exposure risk are available, and suggests a sensitivity analysis when external data are unavailable. Section CECE in Time-to-event Settings extends the results to time-to-event outcomes, in a setting in which individuals can be censored due to loss to follow-up. Section Estimation and Implementation describes how our new parameters can be estimated using existing estimators, even when the outcome is unmeasured. Section Example: Effects of COVID-19 Vaccination implements the new results in a study of the ChAdOx1 nCoV-19 (Oxford) vaccine against COVID-19.

DATA STRUCTURE

Suppose that we have data from a randomized experiment with n individuals who are assigned a binary treatment $A \in \{0, 1\}$ at baseline, where $A = 1$ indicates receiving vaccine and $A = 0$ indicates placebo or other control. As is common in vaccine trial settings,^{16,17} we consider inference in a much larger population from which the trial participants are drawn, so that interactions among patients in the trial are negligible; thus, we suppose the individuals are iid and omit the i subscript. Let L be a vector of baseline covariates. To simplify the presentation, we suppose L is discrete, but the results generalize to continuous L .

Let $E \in \{0, 1\}$ be an indicator of whether an individual is exposed to the infectious agent at least once, for example, being in close contact with an actively contagious individual, which may be *unobserved* in the study. Although we will focus on settings where the exposure E occurs after treatment A is assigned, that is, after baseline, our results also allow the (unobserved) exposure E to be temporally ordered, and thus occur, before treatment A . We first consider $Y \in \mathbb{R}_{\geq 0}$ to be the outcome of interest, for example, disease severity or hospitalization, measured at a given time after randomization, where we define $Y = 0$ when an individual does not have the outcome. In Section CECE in Time-to-event Settings (CECE=causal effect conditional on exposure), we extend the results to censored time-to-event outcomes.

We use superscripts to denote counterfactuals.^{18,19} For example, $Y^{a=1}$ and $Y^{a=0}$ are the outcomes of interest when the treatment is, possibly contrary to fact, fixed to active vaccine ($a = 1$) or control ($a = 0$).

CAUSAL PARAMETERS

The Average Treatment Effect

To motivate the new contributions in this article, we first review the conventional average treatment effect (ATE) of A on the outcome Y ,

$$\mathbb{E}(Y^{a=1}) \text{ vs. } \mathbb{E}(Y^{a=0}), \quad (1)$$

which compares the average outcome in the trial population had everyone been treated ($a = 1$) versus not treated ($a = 0$). This contrast can be identified without additional assumptions when the trial is perfectly executed, that is, under perfect randomization and no losses to follow-up. However, as with any trial, the magnitude of (1) depends on the specific setting in which the RCT was conducted; in a vaccine trial, crucial characteristics include the number of currently infected in the population, the number of previously infected, the mixing pattern, and additional public health measures that may be simultaneously implemented. To generalize the results from the RCT to a policy-relevant setting, we must account for these characteristics, which is far from straightforward.

Conditional Counterfactual Contrasts Are Not Necessarily Causal Effects

To mitigate some of the concerns that are raised about the ATE in vaccine trials, we could attempt to adjust for exposure to the infectious agent.^{2,11} However, defining causal effects conditional on exposure status is not straightforward because exposure status is a post-treatment variable. In particular, a naive contrast of counterfactual outcomes conditional on exposure status,

$$\mathbb{E}(Y^{a=1} \mid E^{a=1} = 1) \text{ vs. } \mathbb{E}(Y^{a=0} \mid E^{a=0} = 1), \quad (2)$$

is not a causal effect when the treatment affects the post-treatment event; it compares counterfactual outcomes in different subpopulations of individuals. This is illustrated by the path $A \rightarrow E$ in the causal directed acyclic graph (DAG) in Figure 1A, which leads to an indirect effect of vaccination on the outcome Y through the path $A \rightarrow E \rightarrow Y$. This indirect effect is plausible if participants know their treatment status; for example, one may expect that vaccinated individuals show a reduction in protective behaviours, which increases the risk of being exposed.

The Principal Stratum Effect and the Causal Effect Conditional on Exposure

A principal stratum effect (PSE)^{18,20} compares counterfactual outcomes among individuals with the same counterfactual exposure status. We can define a particular PSE among those individuals who would be exposed to the infectious agent, at least once, regardless of treatment assignment,

$$\mathbb{E}(Y^{a=1} \mid E^{a=0} = E^{a=1} = 1) \text{ vs. } \mathbb{E}(Y^{a=0} \mid E^{a=0} = E^{a=1} = 1). \quad (3)$$

Unlike (2), the PSE (3) is a contrast of counterfactual outcomes in the same (sub)population of individuals, and it is therefore a causal effect. However, the conditioning set in (3) is defined by exposures in the same individual under two different treatments and, without further assumptions, it is impossible to observe the individuals in this subpopulation,¹⁸ even when E is measured. Thus, the PSE is defined in an




Figure 1. The DAG in (A) describes a study where A is randomly assigned. The DAG in (B) further encodes the no effect on exposure assumption, which is supposed to hold in a *blinded* RCT. The graph in (C) is a SWIG where we have fixed the treatment to a . This SWIG can be used to study identifiability conditions for the CECE, which is identified even if L is unmeasured. The SWIG in (D) describes interventions on both A and E (e is fixed to 1), which allows us to study identifiability conditions for the CDE. Unlike the CECE, the CDE would require measurement of L .

unknown subpopulation that is unobservable even in principle, and the practical relevance of the PSE has been seriously questioned.^{21–24}

As an alternative to the PSE, consider a contrast of counterfactual outcomes conditional on exposure status in the observed data,

$$\mathbb{E}(Y^{a=1} | E = 1) \text{ vs. } \mathbb{E}(Y^{a=0} | E = 1). \quad (4)$$

Like (3), the contrast in (4) is a causal effect as it compares the same subpopulation of individuals under different treatment. Unlike (3), the conditioning set in (4) is observable when E is measured. Without additional assumptions, however, the interpretation of (4) is not straightforward, because an individual’s exposure status in the observed world (E) is not guaranteed to be equal to the exposure status under an intervention that fixes the treatment to be a (E^a). Thus, in general we cannot interpret (4) as a direct effect of treatment A on the outcome Y outside of the treatment effects on exposure status.

But there is at least one setting in which differences in exposure status would not be expected between treatment groups: a *blinded* RCT, which is the context of many vaccine efficacy studies. The following mechanistic assumption formalizes the notion that receiving the vaccine does not exert effects on exposure status E .

Assumption (No effect on exposure).

$$E^{a=0} = E^{a=1}. \quad (5)$$

Assumption (5) guarantees that exposure to the infectious agent is the same, regardless of the treatment that was assigned, and, assuming that the intervention on A is well-defined, allows us to write $E = E^{a=0} = E^{a=1}$.

This assumption can also hold outside of a *blinded* RCT. In particular, exposures that are outside of the individual’s control can satisfy Assumption (5). Such exposures could be consequences of natural or human disasters, such as a flooding after an intense rainfall or radiation from an atomic bombing.

The DAG in Figure 1B describes the causal structure of a *blinded* RCT, in which this assumption would be expected to be met, as there is no path $A \rightarrow E$ and therefore no indirect effect of vaccination on the outcome through the path $A \rightarrow E \rightarrow Y$.

Under assumption (5), the contrasts (2)–(4) are equal, that is,

$$\begin{aligned} & \mathbb{E}(Y^{a=1} | E = 1) \text{ vs. } \mathbb{E}(Y^{a=0} | E = 1) \\ &= \mathbb{E}(Y^{a=1} | E^{a=1} = 1) \text{ vs. } \mathbb{E}(Y^{a=0} | E^{a=0} = 1) \\ &= \mathbb{E}(Y^{a=1} | E^{a=0} = E^{a=1} = 1) \text{ vs. } \mathbb{E}(Y^{a=0} | E^{a=0} = E^{a=1} = 1). \end{aligned}$$

Halloran and Struchiner¹¹ also advocated contrasts of (counterfactual) outcomes in exposed individuals, under the

assumption that “people did not change their behavior after randomization”¹¹ (p. 147). Condition (5) formalizes when such contrasts are unambiguous causal effects, that is, contrast of outcomes in the same (sub)population of individuals.

Because we focus on blinded RCTs in this study, we will use assumption (5) extensively, and under (5), we will denote the contrasts (3)–(4) collectively as the causal effect conditional on exposure (CECE), which is also equal to (2).

The CECE could mitigate some of the concerns that are raised about the generalizability of the ATE (1), because the CECE is confined to those individuals who are exposed to the infectious agent in the observed data, regardless of treatment assignment. Thus, assumption (5) ensures that the CECE has a mechanistic interpretation as an average causal effect given exposure to the infectious agent. The CECE is also of immediate interest for individuals who, based on their own subject-matter knowledge, believe, or possibly know, that they will be, or already are, exposed.

However, the CECE is defined among those who would be exposed in a given study, and the subset who is exposed is context-dependent. To understand the CECE, it is helpful to draw an analogy to ring vaccination trials, in which individuals are recruited only if they have been exposed to an index case, and are subsequently randomly assigned to A . Suppose we indicate exposure to an index case as E . Then, the estimand (4) corresponds to the usual estimand in ring vaccination trials, which is an effect conditional on being exposed. Like the CECE, the target population of a ring vaccination trial is context-dependent, as being a contact of an index case is required for inclusion, and characteristics of those individuals depend on the setting. In a ring vaccination trial, however, exposure is pre-treatment and exposure status is known, features not shared with our setting.

The Controlled Direct Effect

A special case of a controlled direct effect (CDE),²⁵ also called a per-exposure effect or a challenge effect,^{10,11} is defined with respect to an intervention on the treatment A and the exposure E ,

$$\mathbb{E}(Y^{a=1,e=1}) \text{ vs. } \mathbb{E}(Y^{a=0,e=1}). \tag{6}$$

This CDE corresponds to the effect that is identified by a challenge trial²⁶; that is, a study where the participants are subject to an intervention where they are guaranteed to be physically exposed to the infectious agent. Outside of RCTs, household studies are sometimes used to infer such effects, based on contrasts of household secondary attack rates.¹¹

Unlike the ATE (1), the CDE is defined in a controlled setting, in which *all* individuals are exposed to the infectious agent. Thus, this effect is insensitive to the risk of exposure in the observed population.

Finally, all the estimands considered so far can be defined conditional on any baseline covariate L . The distinction

between estimands conditional on L and marginal estimands will be of interest when we study identification.

The Notion of a “Biological” Effect

Both the CECE and CDE quantify treatment effects in individual who are guaranteed to be exposed to the disease agent. In that sense, both effects seem to be captured by the notion of “biological” effects. However, the fact that the CECE and CDE are distinct estimands illustrates that the term “biological” effect, without further clarification, is ambiguous.

IDENTIFICATION

To motivate the identification results in this work, we first review three standard identifiability conditions for the ATE.

Assumption (Treatment exchangeability).

$$Y^a, E^a \perp A \forall a \in \{0, 1\}. \tag{7}$$

Treatment exchangeability, for example, holds in the Single World Intervention Graph (SWIG)¹⁹ in Figure 1C, even if L is unmeasured.

Assumption (Positivity).

$$P(A = a) > 0 \quad \forall a \in \{0, 1\}. \tag{8}$$

Assumption (Consistency).

$$\text{If } A = a, \text{ then } E = E^a \text{ and } Y = Y^a \quad \forall a \in \{0, 1\}. \tag{9}$$

Conditions (7)–(9) hold by design in an RCT where treatment is unconditionally randomly assigned. These three conditions allow us to identify the ATE (1) as $\mathbb{E}(Y | A = 1)$ vs. $\mathbb{E}(Y | A = 0)$, regardless of whether exposure status E is measured.

However, our focus is on estimand (4) (and (6) in eAppendix D, <http://links.lww.com/EDE/B994>), which is defined with respect to counterfactual statuses of the exposure E , so which require additional assumptions.

Identification of the CECE

Under the no effect on exposure assumption (5) and conditions (7)–(9), it is straightforward to express the CECE as a function of factual variables,

$$\begin{aligned} & \mathbb{E}(Y^{a=1} | E = 1) \text{ vs. } \mathbb{E}(Y^{a=0} | E = 1) \\ & = \mathbb{E}(Y | E = 1, A = 1) \text{ vs. } \mathbb{E}(Y | E = 1, A = 0), \end{aligned}$$

but the CECE, as defined as an arbitrary contrast (“vs.”), is not point identified in our data because $\mathbb{E}(Y | E = 1, A = a)$ is not estimable when E is unmeasured. For example, the absolute CECE,

$$\begin{aligned} & \mathbb{E}(Y^{a=1} | E = 1) - \mathbb{E}(Y^{a=0} | E = 1) = \\ & \mathbb{E}(Y | E = 1, A = 1) - \mathbb{E}(Y | E = 1, A = 0), \end{aligned}$$

is not possible to estimate from the observed data.

To identify the CECE, we therefore introduce an additional assumption, which relates the unmeasured E to Y .

Assumption (Exposure necessity).

$$E^a = 0 \implies Y^a = 0, \quad \forall a \in \{0, 1\}. \tag{10}$$

The exposure necessity assumption states that only individuals who were exposed to the infectious agent can experience the outcome. Thus, the exposure is a necessary condition for

experiencing the outcome. For example, contact with some amount of live virus is necessary to develop severe disease. Many exposures and outcomes of interest meet this criterion, though sometimes researchers may be interested in other exposures that do not necessarily satisfy this criterion, for example, sharing a home or classroom with an infected individual. However, such an exposure definition might be revised to being in the same room with an infected individual for at least 1 minute, though even that might not be strictly necessary. In practice, it is important that the investigator has articulated a well-defined exposure, but it is possible that different investigators use different definitions.

Our first proposition shows that the *relative CECE* is identified under the conditions we have introduced so far, which are expected to hold in a blinded RCT.

Proposition 1 (Relative CECE). *Under the no effect on exposure assumption (5), standard identifiability conditions (7)–(9) and exposure necessity (10), the relative CECE is equal to*

$$\frac{\mathbb{E}(Y^{a=1} | E = 1)}{\mathbb{E}(Y^{a=0} | E = 1)} = \frac{\mathbb{E}(Y | A = 1)}{\mathbb{E}(Y | A = 0)},$$

given that $\mathbb{E}(Y | A = 0) > 0$.

The proof is found in eAppendix A, <http://links.lww.com/EDE/B994>. From our considerations in Section The Principal Stratum Effect and the Causal Effect Conditional on Exposure and our derivations in Section Identification of the CECE, it follows that Proposition 1 also gives an identification result for the relative principal stratum effect, that is, $\frac{\mathbb{E}(Y^{a=1} | E^{a=0} = E^{a=1} = 1)}{\mathbb{E}(Y^{a=0} | E^{a=0} = E^{a=1} = 1)}$. Interestingly, Proposition 1 shows that the relative CECE is equal to the conventional ATE on the relative risk scale, which is routinely reported in RCTs. Thus, we specify the assumptions that allow for interpretation of this estimand as a measure of vaccine efficacy conditional on exposure to infection.²

The fact that the relative CECE is identified by the same formula as the relative ATE is related to the known result in epidemiology that diagnostic tests that have perfect specificity will give unbiased estimates of risk ratios, even if these tests do misclassify disease cases. We discuss this in eAppendix C, <http://links.lww.com/EDE/B994>.

Although the absolute CECE is not point identified, our next proposition gives partial identification of the absolute CECE for a binary outcome $Y \in [0, 1]$ in terms of sharp bounds. To simplify the presentation of the subsequent results we suppose, without loss of generality, that $\mathbb{E}(Y | A = 0) \geq \mathbb{E}(Y | A = 1)$.

Proposition 2 (Absolute CECE). *Under the no effect on exposure assumption (5) and conditions (7)–(10), the absolute CECE on an outcome $Y \in [0, 1]$ is partially identified by the sharp bounds*

$$\mathbb{E}(Y | A = 0) - \mathbb{E}(Y | A = 1) \leq \mathbb{E}(Y^{a=0} | E = 1) - \mathbb{E}(Y^{a=1} | E = 1) \leq 1 - \frac{\mathbb{E}(Y | A = 1)}{\mathbb{E}(Y | A = 0)},$$

when $\mathbb{E}(Y | A = 0) \geq \mathbb{E}(Y | A = 1)$.

The proof is given in eAppendix A, <http://links.lww.com/EDE/B994>.

Remark on Proposition 2. The lower bound on the absolute CECE is equal to the absolute ATE. Thus, Proposition 2 gives us another interpretation of a standard risk difference—as a lower bound on the absolute CECE. Furthermore, this lower bound is equal to the absolute CECE if *everybody* is exposed.

The upper bound is 1 minus the relative ATE, which is a quantity that is often reported as the vaccine efficacy in randomized controlled trials,² for example, during the COVID-19 pandemic.²⁷ The absolute CECE is equal to this bound if an unvaccinated individual ($A = 0$) will experience the outcome ($Y = 1$) if and only if she is exposed ($E = 1$).

It follows from Proposition 2 that the larger $\mathbb{E}(Y | A = 0)$, the more informative are the bounds. In particular, the lower bound is equal to the upper bound when $\mathbb{E}(Y | A = 0) = 1$.

Zhao et al.²⁸ studied another interesting setting where relative—but not absolute—risks could be point identified. Their causal question, which concerned racial discrimination in policing, was studied in a setting where the treatment, equivalent to our A , was unmeasured, but the mediator, equivalent to our E , was measured. Their estimand of interest was the conventional ATE.

EXTERNAL DATA AND SENSITIVITY ANALYSIS

Consider a binary outcome $Y \in \{0, 1\}$, for example, an indicator of symptomatic disease. Suppose that the investigator has external knowledge about the risk of experiencing the outcome given exposure among the unvaccinated, that is, $P(Y = 1 | E = 1, A = 0)$. Alternatively, suppose that the investigator has external knowledge about the risk of being exposed among the unvaccinated, that is, $P(E = 1 | A = 0)$. Knowledge of either of these probabilities could have been collected among trial eligible individuals who did not participate in the randomized experiment, or among a subset of the trial participants.

Our next proposition shows that knowledge of either $P(Y = 1 | E = 1, A = 0)$ or $P(E = 1 | A = 0)$ allows point identification of the absolute CECE, when we also assume the same identification conditions as in Proposition 2.

Proposition 3 (Point identification of the absolute CECE). *Under the no effect on exposure assumption (5) and conditions (7)–(10),*

$$\begin{aligned} & \mathbb{E}(Y^{a=0} | E = 1) - \mathbb{E}(Y^{a=1} | E = 1) \\ &= \mathbb{E}(Y | E = 1, A = 0) \left(1 - \frac{\mathbb{E}(Y | A = 1)}{\mathbb{E}(Y | A = 0)} \right) \end{aligned} \quad (11)$$

$$= \frac{\mathbb{E}(Y | A = 0)}{P(E = 1 | A = 0)} - \frac{\mathbb{E}(Y | A = 1)}{P(E = 1 | A = 1)}. \quad (12)$$

The proof of Proposition 3 is given in eAppendix E, <http://links.lww.com/EDE/B994>. Besides giving point identification results in settings with knowledge from external data, Proposition 3 motivates sensitivity analyses for the magnitude of the absolute CECE using sensitivity parameters that are justified by subject-matter reasoning; that is, the investigator can evaluate (11) and (12) under different values of the

marginal sensitivity parameters $P(Y = 1|E = 1, A = 0)$ and $P(E = 1|A = 0)$, respectively.

CECE IN TIME-TO-EVENT SETTINGS

In both RCTs and observational studies, it is common to evaluate vaccine effects on time-to-event outcomes. Our results generalize to settings where the exposure status and the outcome of interest are both time-to-event variables, which possibly are censored due to losses to follow-up.

Suppose that Y_k and E_k are time-to-event variables indicating whether an individual has experienced the event by time k , that is, $Y_k = 1$, and has been exposed by time k , respectively. That is, $E_k = 1$ means exposure has occurred at least once. Let C_k indicate loss to follow-up (censoring) by time $k > 0$. To align with the established causal inference literature,^{18,19,29} suppose that we are interested in outcomes in discrete time intervals $k = 0, \dots, K$, and define the temporal (and topological) order (C_k, E_k, Y_k) in each interval $k > 0$. This setting will converge to a continuous time setting when we let the time intervals become small. We continue to use superscripts to denote counterfactuals, and we formally consider a counterfactual estimand under interventions on the baseline treatment A and the censoring variable C_k .^{30,31} For example, $Y_k^{a,c=0}$ is the counterfactual outcome of interest by time k when treatment is assigned to a and there is no loss to follow-up. The Single World Intervention Graph (SWIG) in Figure 2 describes a causal structure that is coherent with our time-to-event setting. In eAppendix B, <http://links.lww.com/EDE/B994>, we give more details on the time-to-event notation, and we state generalizations of the exchangeability, positivity, consistency, exposure necessity and the no effect on exposure conditions to settings with time-to-event outcomes, see conditions (S3)–(S9). Under these conditions, we can identify the relative CECE as a ratio of cumulative incidences, as described in the next proposition.

Proposition 4 (Relative and absolute CECE for time to event outcomes). *Under exchangeability, positivity, consistency, exposure necessity and the no effect on exposure assumption for time-to-event outcomes, formally stated as conditions (S3)–(S9) in eAppendix B, <http://links.lww.com/EDE/B994>, the relative CECE at time k , $0 \leq k \leq K$, is identified by the ratio of cumulative incidences,*

$$\frac{\mathbb{E}(Y_k^{a=1,c=0} | E_k^{a=1,c=0} = 1)}{\mathbb{E}(Y_k^{a=0,c=0} | E_k^{a=0,c=0} = 1)} = \frac{\mu_k(1)}{\mu_k(0)},$$

where

$$\mu_k(a) = \sum_{s=1}^k h_s(a) \prod_{j=0}^{s-1} [1 - h_j(a)]$$

and

$$h_k(a) = \frac{\mathbb{E}[Y_k(1 - Y_{k-1})(1 - C_k) | A = a]}{\mathbb{E}[(1 - Y_{k-1})(1 - C_k) | A = a]}.$$

Under the same conditions, the absolute CECE is partially identified by the sharp bounds

$$\mu_k(0) - \mu_k(1) \leq \mathbb{E}(Y_k^{a=0,c=0} | E_k^{a=0,c=0} = 1) - \mathbb{E}(Y_k^{a=1,c=0} | E_k^{a=1,c=0} = 1) \leq 1 - \frac{\mu_k(1)}{\mu_k(0)},$$

when $\mu_k(0) \geq \mu_k(1)$.

See eAppendix B, <http://links.lww.com/EDE/B994>, for details and a proof. Thus, like the point exposure and point outcome setting, we do not need to measure common causes of E_j and Y_k , $j, k \in \{0, \dots, K\}$, to point identify the relative CECE and bound the absolute CECE in time-to-event settings.

We have restricted all our discussion to results on *risks*, not *rates* such as hazards. Despite the fact that hazards are sometimes reported as “efficacy parameters” in infectious disease settings, there are well-known limitations of considering causal estimands on the hazard scale^{18,32–34} because of the conditioning on a post-treatment event—here outcomes at earlier times—that is affected by treatment.

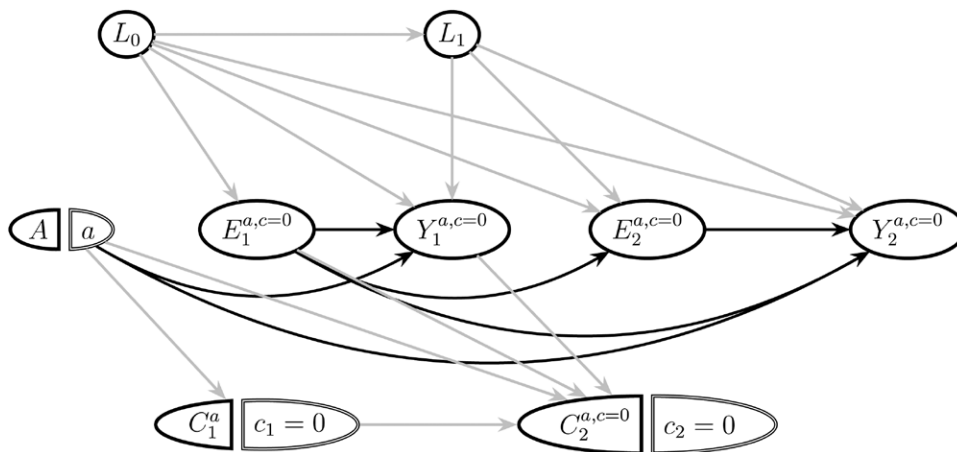


Figure 2. The SWIG shows a time-to-event setting where the CECE is identified, even if L_0 and L_1 are unmeasured.

Excess and Etiologic Fractions

Following Greenland and Robins,³⁵ the excess (prevented) fraction quantifies the excess of outcomes under treatment versus control. When the assumptions of Proposition 4 hold, the excess fraction among the exposed is

$$\frac{\mathbb{E}(Y_k^{a=0,c=0} | E_k^{a=0,c=0} = 1) - \mathbb{E}(Y_k^{a=1,c=0} | E_k^{a=1,c=0} = 1)}{\mathbb{E}(Y_k^{a=0,c=0} | E_k^{a=0,c=0} = 1)} = 1 - \frac{\mu_k(1)}{\mu_k(0)}, \quad (13)$$

which quantifies the increase in caseload under no treatment.^{35,36} In particular, the excess fraction conditional on exposure is equal to the unconditional excess fraction. Furthermore, the right hand side of (13) is often what is reported as *the vaccine efficacy* in clinical studies.²

The excess fraction should not be confused with the etiologic fraction, which is the fraction *caused* by treatment. For example, suppose we consider outcomes at time k , and there are no losses to follow-up. Consider an individual for whom a vaccine prolonged the time to the outcome of severe infection from time j to time l , but (s)he would nevertheless have a severe infection by time k when taking the vaccine, where $j < l < k$. Then, treatment A was a contributory cause of the outcome in this individual, and would thus count as an etiologic event in the etiologic fraction. On the other hand, the individual would not increase the excess caseload by time k , because (s)he experienced the outcome by time k regardless of treatment. The etiologic fraction requires much stronger conditions for identification, even in RCTs and even without conditioning on exposure.^{35,36}

ESTIMATION AND IMPLEMENTATION

Because all our identifying formulas from Section Identification are expressed in terms of simple conditional means, we can use simple estimators with known properties. Let $\hat{\mu}(a)$ and $\hat{\mu}(a, l)$ be estimators of $\mathbb{E}(Y | A = a)$ and $\mathbb{E}(Y | A = a, L = l)$, respectively, for example, empirical means. We can estimate the relative CECE by

$$\widehat{\text{rCECE}} = \frac{\hat{\mu}(1)}{\hat{\mu}(0)},$$

and similarly the upper bound on the absolute CECE by $\widehat{\text{aCECE}}_U = 1 - \widehat{\text{rCECE}}$, where we can compute confidence intervals using standard estimators for risk ratios. Estimators of confidence intervals for risk ratios could be derived from Fieller's theorem³⁷ or the Delta method.³⁸ The estimator of the lower bound on the absolute CECE is $\widehat{\text{aCECE}}_L = \hat{\mu}(0) - \hat{\mu}(1)$, which is simply a difference in means estimator. The estimator for the relative conditional CDE is defined analogously to $\widehat{\text{rCECE}}$, where we also include L in the conditioning set, that is, $\widehat{\text{rCDE}}(l) = \hat{\mu}(1, l) / \hat{\mu}(0, l)$.

For the identifying formulas in Section CECE in Time-to-event Settings, which are cumulative incidences, let $\hat{\mu}_k(a)$ and $\hat{\mu}_k(a, l)$ be estimators of $\mu_k(a)$ and $\mu_k(a, l)$, respectively. Then, we can follow standard procedures for calculating ratios of cumulative incidence functions with confidence intervals, see for example, 39 (Sections 2.3 and 2.4).

EXAMPLE: EFFECTS OF COVID-19 VACCINATION

To study the effect of the ChAdOx1 nCoV-19 vaccine against COVID-19, Voysey et al.⁴⁰ enrolled 23,848 participants in a blinded RCT done across the United Kingdom, Brazil, and South Africa. The participants were randomly assigned 1:1 to the ChAdOx1 nCoV-19 vaccine or control, which contained a meningococcal vaccine. The interim analysis included 11,636 participants.⁴⁰ The cumulative incidence of COVID-19 80 days since second dose was 0.9% (95% CI = 0.5%, 1.3%) and 3.1% (95% CI = 2.4%, 3.8%) in the vaccine and placebo arms, respectively. Thus, an estimate of the relative CECE $\equiv \text{CECE}_{k=80}$, defined on the cumulative incidence scale, is

$$\widehat{\text{rCECE}} = \frac{\hat{\mu}(1)}{\hat{\mu}(0)} = 0.30 \quad (95\% \text{ CI} = 0.15, 0.44),$$

which corresponds to the reported vaccine efficacy point estimate of $1 - 0.30 = 0.70$ ⁴⁰ (Table 2). Here and henceforth, we omit the $k = 80$ subscript to simplify the notation.

We can use the results from Section Identification of the CECE to derive bounds for the absolute CECE, specifically the sharp lower bound

$$\widehat{\text{aCECE}}_L = 0.031 - 0.009 = 0.022 \quad (95\% \text{ CI} = 0.011, 0.033),$$

and the sharp upper bound

$$\widehat{\text{aCECE}}_U = 1 - 0.30 = 0.70 \quad (95\% \text{ CI} : 0.57 - 0.85).$$

Although we obtained informative point estimates of the relative CECE, the bounds on the absolute CECE are wide. The fact that the bounds are wide is not surprising, because they crucially depend on the risk of exposure to the virus. As discussed in Section Identification of the CECE, the lower bound is reached under a setting where everybody is exposed to the virus, and the upper bound when the probability of the outcome among the exposed, unvaccinated individuals is 1. Depending on the definition of exposure, such settings may or may not be plausible. However, we can use a sensitivity analysis, as suggested in Section External Data and Sensitivity Analysis, to reason about the magnitude of the absolute CECE.

Sensitivity Analysis in the ChAdOx1 nCoV-19 Vaccine Study

Determining sensitivity parameters to generate point estimates of the absolute CECE requires us to think concretely about the definition of exposure, or to consider a range of definitions of exposure. The ChAdOx1 nCoV-19 vaccine trial began enrollment in June 2020 and recruited a sample of 60%–90% health-care workers, depending on the site. Suppose we define E as coming into contact with an equivalent amount of SARS-CoV-2 virus particles that may be encountered while caring for a COVID patient wearing personal protective equipment (PPE). However, because it is important that our definition of E satisfies exposure necessity, we can more precisely define E as a specific amount of virus particles such that the

exposure necessity condition holds. For example, this could be a particular amount of virus particles when wearing PPE, and a higher amount when not wearing PPE. To parameterize $P(E = 1|A = 0)$, we might propose that 60% of the trial participants were exposed to such an amount of virus particles at some point during the 80-days period. Given the observed data, this would imply that $P(Y = 1|E = 1, A = 0) = 0.052$; that is, E , here denoting a given amount of virus particles, was sufficient to cause symptomatic COVID-19 in just over 5% of unvaccinated participants during 80 days of follow-up. In this setting, we would estimate $\widehat{aCECE} = 0.037$ (Figure 3). Suppose now that we rather set the sensitivity parameter $P(E = 1|A = 0)$ to 0.9 instead of 0.6. Then, $\widehat{aCECE} = 0.024$, which is consistent with $P(Y = 1|E = 1, A = 0) = 0.034$.

So far we have reasoned about the sensitivity parameter $P(E = 1|A = 0)$. However, we could also reason about $P(Y = 1|E = 1, A = 0)$, perhaps using external data. For example, consider the choir practice in Washington state in March 2020, after which 52 out of 61 participants developed COVID-19, having been exposed to a high concentration of virus in an unmasked setting.⁴¹ Using this estimate of $P(Y = 1|E = 1, A = 0) = 0.85$, we find that $\widehat{aCECE} = 0.60$, an estimate consistent with $P(E = 1|A = 0) = 0.036$. If exposure to such a high dose of SARS-CoV-2 is necessary for infection, the absolute CECE is much closer to its upper bound, and the risk of such an exposure in the trial setting is necessarily lower.




Figure 3. Illustration of the relation between the sensitivity parameters $P(Y = 1|E = 1, A = 0)$ and $P(E = 1|A = 0)$. Specifying one sensitivity parameter is sufficient to get a point estimate of the absolute CECE, as illustrated by three values of \widehat{aCECE} in the figure.

More broadly, the bounds illustrate an important point: the relative CECE is constant for *any* exposure when Assumptions (10) and (5) hold, but only weak conclusions can be made about the magnitude of the absolute CECE unless we both have a clear idea about the definition of the exposure and have information about $P(Y = 1|E = 1, A = 0)$ or $P(E = 1|A = 0)$.

Estimating the CDE requires data on covariates L , such as comorbidities, smoking, work occupation, and age, to justify condition (S10), see eAppendix D, <http://links.lww.com/EDE/B994>. Because this information is unavailable, we have not attempted to estimate the CDE.

DISCUSSION

In this study, we have distinguished various estimands for vaccine effects conditional on exposure to infection and clarified their identification assumptions. We have required that the exposure, for example, close contact with an infectious individual, is necessary for the outcome of interest to occur, for example, symptomatic disease, as stated in our exposure necessity condition (10). An alternative approach would involve adapting the definition of exposure to something that is possible to measure. For example, one might define exposure as close contact with infected people who present overt disease. However, such definitions have explicitly been discouraged, precisely because they would lead to an *underestimate* of the exposure in settings where some infections are inapparent.² In the case of the CECE, we have considered the exposure to be any event such that the exposure necessity condition holds.

When a necessary exposure is unmeasured, we have shown that relative effects can be point identified under plausible conditions, but absolute effects can only be bounded under the same conditions. Often the most commonly reported and publicized results are relative effects, as in major studies on different COVID-19 vaccines.^{40,42,43} Thus, the results presented in this work give valuable interpretations to the numbers that are computed.

However, often both relative and absolute effects are of interest. Absolute effects are usually studied in optimal regime settings,⁴⁴⁻⁴⁶ which reflects the common opinion that heterogeneous effects on the additive scale are most appropriate for evaluating public health interventions.⁴⁷ Importantly, bounds on the additive effect can be used in formal decision theoretic approaches, even if these bounds are wide or cover null effects.^{48,49} Furthermore, if the investigator is willing to invoke assumptions about the probability of exposure, the bounds will be narrower, as we describe in Section External Data and Sensitivity Analysis.

In future study, we will formally consider generalizability of the different vaccine effects on different scales, including the CECE, which could be applicable to settings with interference outside of the randomized experiment.

ACKNOWLEDGMENTS

This study was initiated during Louisa Smith's visit to École Polytechnique Fédérale de Lausanne.

REFERENCES

- Greenwood B. The contribution of vaccination to global health: past, present and future. *Philos Trans R Soc B Biol Sci.* 2014;369:20130433.
- Halloran ME, Struchiner CJ, Longini IM. *Design and Analysis of Vaccine Studies*, vol 18. Springer; 2010.
- Mehrotra DV, Janes HE, Fleming TR, et al. Clinical endpoints for evaluating efficacy in covid-19 vaccine trials. *Ann Intern Med.* 2021;174:221–228.
- Gilbert PB, Fong Y, Carone M. Assessment of immune correlates of protection via controlled vaccine efficacy and controlled risk. *arXiv preprint arXiv:2107.05734.* 2021.
- Lipsitch M, Dean NE. Understanding covid-19 vaccine efficacy. *Science.* 2020;370:763–765.
- Lipsitch M, Kahn R. Interpreting vaccine efficacy trial results for infection and transmission. *Vaccine.* 2021;39:4082–4088.
- Kilpatrick KW, Hudgens MG, Elizabeth Halloran M. Estimands and inference in cluster-randomized vaccine trials. *Pharm Stat.* 2020;19:710–719.
- Follmann DA, Fay MP. Vaccine efficacy at a point in time. *medRxiv.* 2021.
- Patel MK, Bergeri I, Bressee JS, et al. Evaluation of post-introduction covid-19 vaccine effectiveness: Summary of interim guidance of the world health organization. *Vaccine.* 2021;39:4013–4024.
- O'Hagan JJ, Lipsitch M, Hernán MA. Estimating the per-exposure effect of infectious disease interventions. *Epidemiology (Cambridge, Mass.).* 2014;25:134.
- Halloran ME, Struchiner CJ. Causal inference in infectious diseases. *Epidemiology.* 1995;142:151.
- Jamrozik E, Selgelid MJ. Covid-19 human challenge studies: ethical issues. *Lancet Infect Dis.* 2020;20:e198–e203.
- Corey L, Mascola JR, Fauci AS, Collins FS. A strategic approach to covid-19 vaccine r&d. *Science.* 2020;368:948–950.
- Cohen J. Studies that intentionally infect people with disease-causing bugs are on the rise. *Science.* 2016;10.
- Struchiner CJ, Halloran ME. Randomization and baseline transmission in vaccine field trials. *Epidemiol Infect.* 2007;135:181–194.
- Tsiatis AA, Davidian M. Estimating vaccine efficacy over time after a randomized study is unblinded. *arXiv preprint arXiv:2102.13103.* 2021.
- Halloran ME, Longini IM Jr, Struchiner CJ. Estimability and interpretation of vaccine efficacy using frailty mixing models. *Am J Epidemiol.* 1996;144:83–97.
- Robins JM. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Math Model.* 1986;7:1393–1512.
- Richardson TS, Robins JM. Single world intervention graphs (swigs): a unification of the counterfactual and graphical approaches to causality. *Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper.* 2013;128.
- Frangakis CE, Rubin DB. Principal stratification in causal inference. *Biometrics.* 2002;58:21–29.
- Robins JM, Rotnitzky A, Vansteelandt S. Principal stratification designs to estimate input data missing due to death-discussion. *Biometrics.* 2007;63:650–653.
- Joffe M. Principal stratification and attribution prohibition: good ideas taken too far. *Int J Biostatistics.* 2011;7:1–22.
- Dawid P, Didelez V. Imagine a can opener. The magic of principal stratum analysis. *Int J Biostatistics.* 2012;8:19.
- VanderWeele TJ. Principal stratification—uses and limitations. *Int J Biostatistics.* 2011;7:1–14.
- Robins JM, Greenland S. Identifiability and exchangeability for direct and indirect effects. *Epidemiology.* 1992;143:155.
- Hudgens MG, Gilbert PB. Assessing vaccine effects in repeated low-dose challenge experiments. *Biometrics.* 2009;65:1223–1232.
- Kahn R, Wang R, Leavitt SV, Hanage WP, Lipsitch M. Leveraging pathogen sequence and contact tracing data to enhance vaccine trials in emerging epidemics. *Epidemiology.* 2021;32:698–704.
- Zhao Q, Keele LJ, Small DS, Joffe MM. A note on post-treatment selection in studying racial discrimination in policing. *arXiv preprint arXiv:2009.04832.* 2020.
- Hernán MA, Robins JM. *Causal Inference.* CRC; 2018.
- Robins JM, Finkelstein DM. Correcting for noncompliance and dependent censoring in an aids clinical trial with inverse probability of censoring weighted (ipcw) log-rank tests. *Biometrics.* 2000;56:779–788.
- Young JG, Stensrud MJ, Tchetgen Tchetgen EJ, Hernán MA. A causal framework for classical statistical estimands in failure-time settings with competing events. *Stat Med.* 2020;39:1199–1236.
- Hernán MA. The hazards of hazard ratios. *Epidemiology (Cambridge, MA).* 2010;21:13.
- Stensrud MJ, Aalen JM, Aalen OO, Valberg M. Limitations of hazard ratios in clinical trials. *Eur Heart J.* 2019;40:1378–1383.
- Stensrud MJ, Hernán MA. Why test for proportional hazards? *JAMA.* 2020;323:1401–1402.
- Greenland S, Robins JM. Conceptual problems in the definition and interpretation of attributable fractions. *Am J Epidemiol.* 1988;128:1185–1197.
- Robins JM, Greenland S. Estimability and estimation of excess and etiologic fractions. *Stat Med.* 1989;8:845–859.
- Fieller EC. Some problems in interval estimation. *J R Stat Soc Series B (Methodological).* 1954;16:175–185.
- Herson J. Fieller's theorem vs. the delta method for significance intervals for ratios. *J Stat Comput Simul.* 1975;3:265–274.
- Zhang M-J, Fine J. Summarizing differences in cumulative incidence functions. *Stat Med.* 2008;27:4939–4949.
- Vovsey M, Clemens SAC, Madhi SA, et al. Safety and efficacy of the chadox1 ncov-19 vaccine (azd1222) against sars-cov-2: an interim analysis of four randomised controlled trials in brazil, south africa, and the uk. *Lancet.* 2021;397:99–111.
- Hamner L. High sars-cov-2 attack rate following exposure at a choir practice—skagit county, washington, march 2020. *MMWR Morb Mortal Wkly Rep.* 2020;69:606–610.
- Baden LR, El Sahly HM, Essink B, et al. Efficacy and safety of the mrna-1273 sars-cov-2 vaccine. *N Engl J Med.* 2021;384:403–416.
- Thomas SJ, Moreira ED Jr, Kitchin N, et al. Safety and efficacy of the bnt162b2 mrna covid-19 vaccine through 6 months. *N Engl J Med.* 2021;385:1761–1773.
- Murphy SA. Optimal dynamic treatment regimes. *J R Stat Soc Series B (Statistical Methodology).* 2003;65:331–355.
- Robins JM. Optimal structural nested models for optimal sequential decisions. In *Proceedings of the second seattle Symposium in Biostatistics.* Springer, 2004, p. 189–326.
- Tsiatis AA, Davidian M, Holloway ST, Laber EB. *Dynamic Treatment Regimes: Statistical Methods for Precision Medicine.* Chapman and Hall/CRC; 2019.
- Knol MJ, VanderWeele TJ. Recommendations for presenting analyses of effect modification and interaction. *Int J Epidemiol.* 2012;41:514–520.
- Cui Y. Individualized decision-making under partial identification: three perspectives, two optimality results, and one paradox. *Harvard Data Sci Rev.* 2021;27:1397–1421.
- Manski CF. Reasonable patient care under uncertainty. *Health Econ.* 2018;27:1397–1421.