

The Estimands Framework and Causal Inference: Complementary not Competing Paradigms

Thomas Drury^{1,*}, Jonathan W. Bartlett², David Wright³, and Oliver N. Keene⁴

¹GSK, London, UK

²The London School of Hygiene and Tropical Medicine, London, UK

³Statistical Innovation, AstraZeneca, Cambridge, UK

⁴KeenONStatistics, London, UK

*Email: thomas.a.drury@gsk.com

Published in Pharmaceutical Statistics 23 August 2025: <https://doi.org/10.1002/pst.70035>

Abstract: The creation of the ICH E9 (R1) estimands framework has led to more precise specification of the treatment effects of interest in the design and statistical analysis of clinical trials. However, it is unclear how the new framework relates to causal inference, as both approaches appear to define what is being estimated and have a quantity labelled an estimand. Using illustrative examples, we show that both approaches can be used to define a population-based summary of an effect on an outcome for a specified population and highlight the similarities and differences between these approaches. We demonstrate that the ICH E9 (R1) estimand framework offers a descriptive, structured approach that is more accessible to non-mathematicians, facilitating clearer communication of trial objectives and results. We then contrast this with the causal inference framework, which provides a mathematically precise definition of an estimand, and allows the explicit articulation of assumptions through tools such as causal graphs. Despite these differences, the two paradigms should be viewed as complementary rather than competing. The combined use of both approaches enhances the ability to communicate what is being estimated. We encourage those familiar with one framework to appreciate the concepts of the other to strengthen the robustness and clarity of clinical trial design, analysis, and interpretation.

Keywords: ICH E9(R1), Estimands, Causal Inference, Potential Outcomes

1 Introduction

There has been widespread adoption of the estimands framework by the pharmaceutical industry following the publication of ICH E9 (R1)^[1–3]. There has also been some use of the framework in academia^[4]. The purpose of the framework is to provide context and clarity on what is being estimated in a clinical trial by formally defining a quantity labelled an estimand that has five components: Population, Treatments, Endpoint, Summary Measure, and Intercurrent Events (with strategies chosen to address them). The framework was designed to provide a common language for different stakeholders

to discuss and define effects of interest in clinical trials^[5]. It was also, in part, an acknowledgment that to gain a full understanding of an effect targeted in randomized clinical trials (RCTs), we often need to consider how post baseline events impact the outcomes of interest.

Although the framework is a positive step forward in the design, analysis and reporting of clinical trials, a potentially confusing factor is that the term “estimand” has been used to describe a similar idea in the Causal Inference literature for many years^[6]. A further complication is that what is labelled an estimand in each approach appears to partially overlap, and both relate to effects on patient outcomes under different treatment conditions.

Previous work has explored the links between the estimands framework and causal inference^[7–12]. Scharfstein (2019) provided a critique of the draft version of ICH E9 (R1) with high-level links to causal concepts. Lipkovich et al. (2019) used the terminology and principles of causal inference as a framework for defining estimands. Other authors have discussed the linkages to health technology assessments^[11] and targeting estimands with hypothetical IE strategies via instrumental variables^[12].

In this paper we review how ICH E9 (R1) and the causal inference literature define what effect is being estimated in the context of a randomized trial. We discuss the origins of the two approaches, how they relate, and provide examples derived from real trials to illustrate how estimands are defined from both approaches. Although there are pros and cons to both approaches, they share the overall aim of providing clarity on what is being estimated, and we suggest that it can be beneficial to use aspects from both when defining the treatment effect targeted in clinical trials. We also show that, although not explicitly referenced, the ICH E9 (R1) estimands framework relies on concepts developed as part of causal inference methodology. Finally, we discuss our interpretation of some more ambiguous aspects of ICH E9 (R1) which are arguably not clear and comment on some recent opinions in the causal inference literature about the utility of the estimand framework.

This paper is organised as follows: in section 2 we outline some examples of clinical trials describing potential effects that could be targeted, section 3 gives a brief history and outline of the ICH E9 (R1) estimands framework, section 4 reviews relevant aspects of causal inference. Section 5 illustrates how the example effects from section 2 could be formulated using both approaches. Section 6 contrasts and discusses several aspects of both the estimands framework and causal inference and finally we provide a conclusion in section 7.

2 Example Clinical Trials

Randomised clinical trials provide a rigorous basis for comparing groups of patients assigned different pre-specified treatments. With complete follow-up, randomization guarantees that the observed treatment effect on the outcome is, aside from random error, due to the treatment each individual was assigned. However, this effect of assignment is a combination of the effect of taking the assigned treatment plus any subsequent effects from other events that occur after assignment, for instance, treatment discontinuation or use of additional (or rescue) medications. In many clinical trials, this effect of assignment may not be the most relevant clinical question of interest to all trial stakeholders and other questions are often important, including the effects attributable to specific treatments (or combinations of treatments).

Events such as treatment discontinuation or use of additional medication can affect the value and interpretation of the outcome measured. For some clinical questions of interest, this can lead to difficulty using the recorded outcome measures directly to address the research questions in the trial and it is typically necessary to provide more information about how these events are reflected in the effects being targeted.

We present four examples based on real clinical trials that exhibit some of the issues discussed above. We do not discuss the trials in detail but use some aspects as motivation to aid the discussion of the similarities and differences between the causal inference and ICH E9 (R1) approaches to defining what is to be estimated.

2.1 Diabetes

The PIONEER 1 trial compared oral semaglutide to placebo in adult patients with type 2 diabetes, with a primary endpoint of change from baseline HbA1c^[13]. During the trial, patients could discontinue assigned treatment and/or receive rescue medication if they did not maintain acceptable glycaemic control, both of which would likely change their HbA1c. Overall, more placebo patients needed rescue medication compared to semaglutide, creating a differential impact between the groups. Therefore, comparing change in HbA1c values at the end of the trial considers the combined effect of the *semaglutide taken plus any rescue taken* and does not address the treatment effect caused by or attributable to *taking only semaglutide itself for the duration of the trial versus only placebo for the duration of the trial*.

2.2 Nasal Polyps

The SYNAPSE trial compared mepolizumab and placebo in chronic rhinosinusitis with nasal polyps after 52 weeks^[14]. The co-primary endpoints were nasal polyps score and nasal blockage score at the end of the trial. A post baseline event was surgery for nasal polyps and treatment with mepolizumab was intended to reduce the number of patients having to undergo surgery. Although surgery is likely to improve subsequent nasal scores, it represents a failure of randomised treatment. A key question is how to reflect this adverse outcome of nasal surgery in the comparison of treatments. Comparing nasal polyps scores without regard to surgery would ignore this important outcome and not address the causal effect of treatment on the overall clinical outcome.

2.3 COVID19

In a COVID-19 vaccine trial^[15] participants at high risk for SARS-CoV-2 infection or its complications were randomly assigned to receive two intramuscular injections of the mRNA-1273 vaccine or placebo 28 days apart. The primary end point was prevention of COVID-19 illness (defined as onset at least 14 days after the second injection) in participants who had not previously been infected with SARS-CoV-2. Of the 30,351 participants who received at least one dose of either the active vaccine or placebo, 941 participants were excluded from the primary analysis, mainly because of failure to take the second dose or not taking the second dose in the appropriate time window. Clinical interest centred on answering a

causal question of the treatment effect of receiving both doses in the appropriate time window, versus placebo.

2.4 Arthritis

In a safety trial^[16], patients at increased cardiovascular risk that use nonsteroidal anti-inflammatory drugs (NSAIDs) for osteoarthritis or rheumatoid arthritis were randomly assigned to receive celecoxib, ibuprofen, or naproxen. The goal of the trial was to assess the noninferiority of celecoxib with regard to the composite outcome of cardiovascular death, nonfatal myocardial infarction, or nonfatal stroke. Of the 24,081 patients randomized, 68.8% stopped taking their assigned study drug, and 27.4% withdrew from study follow-up. An important question given the safety aims of the trial was the cardiovascular safety of patients while receiving the drugs under study. Given the large percentage of patients that stopped taking treatment, this should be reflected in any estimate of cardiac risk attributed to each treatment.

3 The ICH Estimands Framework

The impetus for developing the estimands framework and revising ICH E9 originated from the seminal 2010 report by the US National Research Council (NRC) on the prevention and treatment of missing data in clinical trials^[17]. The report raised issues with the standard practice at the time for the design and analysis of clinical trials in regard to how missing data were handled. Importantly the report emphasised the role of the trial estimand, stating that “*Estimation of the primary (causal) estimand, with an appropriate estimate of uncertainty, is the main goal of a clinical trial*”. In addition, there was a continuing concern that “per-protocol” type analysis which only includes the patients who complied fully with the trial protocol risked creating groups that were no longer similar and may not reflect a valid causal comparison^[18–21]. As a result of these concerns, regulators led a working group to develop the estimands framework to be included as a revision to ICH E9.

ICH E9 (R1) was finalised in 2019 and we only give an outline for the purposes of our discussions. In brief, the guidance defines a process that starts by determination of the trial objective and associated clinical question of interest. Once this has been agreed, a population-based quantity called “an estimand” (which we label ICH estimand) should be defined to address the clinical question of interest. Once this ICH estimand has been defined, an appropriate estimator can be selected. **The guidance defines an ICH estimand to have five components:**

- Population – the group of patients the treatment(s) are designed to benefit.
- Variable – the outcome measure used to assess the benefit for each patient.
- Treatments – the specific treatment(s) assigned to each group of patients and any comparisons of groups to be made.
- Summary – the population-level summary of the variable of interest.
- Intercurrent Events (IEs) and strategies – the post baseline events that could impact the value or interpretation of the variable measured and how these events are accounted for if they occur.

The framework specifies that **all five components are required to define “the estimand”** and taken together, they provide enough context to fully understand what effect is being estimated. The guidance defines five basic strategies to address IEs. These strategies provide clarity about how the occurrence of IEs are reflected in the estimand being defined. Again, in brief the strategies can be understood as:

1. Treatment Policy – Any effects resulting from the occurrence of IEs are included in the overall effect of interest. If this strategy is applied to all the identified IEs, then this corresponds to the effect of treatment assignment.
2. Hypothetical – Any effects resulting from the occurrence of the IE are removed from the effect of interest. This targets an effect in the absence of these IEs occurring.
3. Composite – The occurrence of the IE is considered an informative but (generally) poor outcome for the patient and the outcome variable’s definition is somehow re-defined so its value changes if the IE occurs.
4. While on Treatment – The focus of interest is the variable measure until the particular IE occurs. The label of this strategy is arguably ambiguous and in general it can be thought of as “while the IE has not occurred”.
5. Principal Stratum – generally, this targets an effect in a group of patients that would (or would not) have had the IE if they had been assigned to one, or either (or any) of the treatments in the trial. This attempts to consider effects in a group of patients where the occurrence (or not) of the IE does not depend on their treatment assignment.

Most ICH estimands defined in later stage trials (Phase IIb and III) will have two or more IEs. The most frequent IEs in general are discontinuation of assigned treatment and/or use of additional (sometimes rescue) medication. Death may be an important IE in trials where it is not the outcome, but depending on the disease may occur infrequently. In trials measuring time to a particular type of event, failures from competing events, or death, may be IEs.

The ICH estimand framework itself has very little mathematical content and relies on text descriptions to communicate concepts and scientific research questions. The text has evolved in two basic forms: a table listing each of the five components explicitly, or a paragraph that incorporates each of the five components^[22]. The details for the IE strategies are also either explicitly stated, or alternatively incorporated into the definitions of treatment, population, or variable attributes. Different opinions exist for which of these approaches provides the most clarity. Some feel strongly that all IEs should be incorporated directly into the other attributes relevant to the handling strategy used. Others suggest that explicitly listing all IEs is clearer. We provide both a table and paragraph description for all ICH Estimands in this paper, but we choose not to incorporate every IE into the other attributes.

ICH E9 (R1) provided minimal guidance on estimation but did emphasise the importance of ensuring any estimand that is defined can be reliably estimated. It also stressed the need to assess the sensitivity of conclusions to statistical assumptions. Regression modelling and missing data approaches such as multiple imputation and inverse probability weighting may be useful but given the complex definitions of what effect an ICH estimand targets, estimation is an evolving field with many opportunities for novel methods^[23–31]. In particular, some estimators from causal inference are being explored^[32].

4 Estimands in Causal Inference

While there are actually multiple causal inference frameworks^[33–35], they all share the aims of providing a language to clearly define causal quantities of interest and to spell out the assumptions under which these quantities would be correctly estimated. Within the causal inference literature, the potential outcomes (counter-factual) framework^[33] is arguably the most popular and is the methodology we use for comparison with the ICH estimands framework. This framework was devised by Rubin and expanded by Hernan and Robins^[6], among others^[35]. Pearl (2022), Robins (2022) and Rubin (2022) have provided a thorough discussion of the history of causal inference^[36–38]. In order to discuss estimands defined in causal inference, we give a brief description of the potential outcomes paradigm.

4.1 Potential Outcomes Framework

We consider the simple setting with individual patients that have a particular illness, and we wish to assess any benefit they might receive from a new drug in comparison to a control. In order to measure the benefit, we use an outcome measure Y and consider the existence of two potential outcomes for Y in the theoretical scenarios where the individual is assigned the control or new drug. For ease of notation, we define the actual treatment assignment variable for an individual as $A = 0, 1$ and a particular (theoretical) treatment assignment as a superscript $a = 0, 1$ above the outcome. We then have two potential outcomes $Y^{(a=0)}$ and $Y^{(a=1)}$ which correspond to the outcomes the individual would experience if they were assigned the control or new drug respectively. The theoretical causal effect of the new drug on the individual could then be considered as the difference in these potential outcomes $Y^{(a=1)} - Y^{(a=0)}$.

The potential outcomes framework relies on imagined worlds where different treatments are assigned to the same individual, but in reality, only one world is observed where the individual is assigned a single treatment. This means we cannot directly estimate individual causal effects. However, under certain conditions, we can estimate average effects in the population of individuals such as the difference in expected potential outcomes for individuals under each treatment $\mathbb{E}[Y^{(a=1)}] - \mathbb{E}[Y^{(a=0)}]$ which is equal to the average of the individual level causal effects. The basic conditions that need to hold in order to estimate this are referred to in the causal inference literature as: *Consistency*, *Positivity* and *Exchangeability*. We only give a very basic idea of these assumptions for the purpose of this work. For a thorough discussion see “What If” by Hernan and Robins^[6]:

- Consistency – this means that the observed outcome under the treatment an individual is actually assigned corresponds to the individual’s potential outcome if assigned that treatment.
- Positivity – this means that there is a non-zero chance that the individual could have received both (or any) of the treatments being compared.
- Exchangeability – this means the chance of a patient being assigned each treatment is independent of the two potential outcomes.

In randomized trials, these assumptions typically hold as randomization ensures a non-zero chance of being independently assigned to each treatment, so in addition to consistency, both positivity and

exchangeability are satisfied. Under these assumptions, the expected values of the potential outcomes are equal to the expected values for the groups of individuals assigned to each treatment in the trial. Specifically, $\mathbb{E}[Y^{(a=1)}] = \mathbb{E}[Y | A = 1]$ (and similarly for control). This means the difference in expected potential outcomes $\mathbb{E}[Y^{(a=1)} - Y^{(a=0)}]$ can be estimated using $\mathbb{E}[Y | A = 1] - \mathbb{E}[Y | A = 0]$. When the treatment effect of interest can be expressed in terms of the observed data quantities, it is said to be *identifiable*.

In observational studies where assignment is not random by design, the positivity and exchangeability assumptions are often not satisfied. The exchangeability assumption can be weakened to a conditional version, which says that treatment assignment is independent of potential outcomes conditional on confounders. In this case the effects are only identifiable if the estimation methods account for all the assumed confounders W using methods such as outcome regression models, or inverse probability weighting with propensity scores.

In more complicated settings, more complex causal ideas, notation, and assumptions are needed. As an example, in trials with longitudinal follow up, a patient’s treatment may change over time and there are different potential treatment sequences they could experience. Expanding our previous set up to include a variable $D_j = 1, 0$ which corresponds to the situation of a patient receiving their treatment correctly at visit j , we create a vector of indicators $\overline{D} = [D_1, \dots, D_J]$ that represent a treatment sequence. For example, with four visits, an imposed sequence of $\overline{d} = [1, 1, 0, 0]$ would indicate a theoretical scenario where a patient receives their assigned treatment for the first two visits, but not the final two visits.

Although repeated follow up increases the complexity, for each potential treatment assignment A and sequence \overline{D} we are still able to define potential outcomes $Y^{(A, \overline{D})}$ for a patient and consider these potential outcomes under specified (sometimes referred to as intervened) treatment assignment and sequences $Y^{(a, \overline{d})}$. These types of treatment sequences are known as so-called “static regimes”. More recent work in causal inference has expanded on this to “dynamic regimes” where the treatment of a patient can depend on time-varying follow up measures^[39].

A tool developed in the causal inference literature to communicate the assumed structural relationships between the outcome, assignment process, and other factors are causal graphs. There are different types of causal graphs, developed for the different causal frameworks, but we consider the most popular which is the directed acyclic graph (DAG)^[40]. In general, these graphs have arrows to indicate assumed causal effects which help practitioners visualize and communicate how different variables affect each other. We include a simple example of a DAG in Figure 1, but for brevity, we do not illustrate DAGs further in this paper.




Figure 1: A simple directed acyclic graphs showing one possible causal structure of a clinical trial where baseline factors W and treatment assignment at randomization A have a causal effect on the outcome Y but also have a causal effect on patient discontinuation from randomized treatment D which also impacts the outcome.

A wide variety of estimation methods have been developed within the causal inference literature, but they can generally be grouped as matching, inverse probability weighting and the so called “G-methods”, which includes G-formula and G-estimation. Combinations of these methods can also be used to create “doubly robust” approaches. All methods usually involve some form of (semi) parametric regression model. For a description of these methods see “What If”^[6].

4.2 Causal Estimands

In the causal inference literature, an estimand (which we call a causal estimand) is a **mathematical expression to represent a summary comparison of potential outcomes under specified treatment conditions for a given population of individuals**. It represents a comparison of quantities that exist for a population and does not depend on which estimation method is used. The most basic example of a causal estimand is the difference in expected values for two treatments: $\mathbb{E}[Y^{(a=1)}] - \mathbb{E}[Y^{(a=0)}]$, but other simple examples include the risk difference for two treatments: $\mathbb{P}(Y^{(a=1)} = 1) - \mathbb{P}(Y^{(a=0)} = 1)$ or the risk ratio for two treatments $\mathbb{P}(Y^{(a=1)} = 1) / \mathbb{P}(Y^{(a=0)} = 1)$. An example of a simple estimand in the repeated follow up setting discussed could be the difference in the expected values of outcome at the final follow-up between two treatments if they were correctly administered across all visits, specifically: $\mathbb{E}[Y_J^{(a=1, \vec{d}=[1,\dots,1])}] - \mathbb{E}[Y_J^{(a=0, \vec{d}=[1,\dots,1])}]$.

5 Examples Revisited

We now revisit the examples discussed in section two and illustrate how these effects of interest could be specified in both the ICH estimand and causal inference frameworks.

5.1 Example 1 - Diabetes

For this example, we only concentrate on the comparison of the 14 mg daily dose of semaglutide and placebo. We consider a situation where the objective of PIONEER 1 might have been to answer these two different clinical questions:

1. What is the effect of being assigned semaglutide vs placebo on HbA1c values after a year.
2. What is the effect of semaglutide vs. placebo if patients were treated continuously for one year.

During the PIONEER 1 trial, a number of patients stopped taking the treatment they were assigned. A number also received rescue medication for glycaemic control. Both of these post baseline events are likely to impact the HbA1c efficacy measure.

5.1.1 Estimands Framework

In the estimands framework discontinuing assigned treatment and initiation of rescue medication would be considered IEs. These would then be included as the fifth component of the ICH estimand together with handling strategies that align with the clinical question of interest above. We provide potential ICH estimands for each question below:

ICH Estimand 1.1

Difference in mean change from baseline HbA1c at week 26 for adult patients with type 2 diabetes assigned to 14 mg daily oral semaglutide or placebo including any subsequent effects of discontinuation of assigned treatment or initiation of rescue medication.

Population	Adult patients with type 2 diabetes
Treatment	14 mg daily oral semaglutide or placebo, including any subsequent discontinuation or initiation of rescue medication
Variable	Change from baseline HbA1c at week 26
Summary Measure	The difference in means
IEs (Handling)	Discontinuation of assigned treatment (Treatment Policy) Initiation of rescue medication (Treatment Policy)

ICH Estimand 1.2

Difference in mean change from baseline HbA1c at week 26 for adult patients with type 2 diabetes between 14 mg daily oral semaglutide and placebo if participants take their assigned treatment for the duration of the treatment period and rescue medication was not available.

Population	Adult patients with type 2 diabetes
Treatment	14 mg daily oral semaglutide or placebo taken throughout the full follow-up period
Variable	Change from baseline HbA1c at week 26
Summary Measure	The difference in means
IEs (Handling)	Discontinuation of assigned treatment (Hypothetical) Initiation of rescue medication (Hypothetical)

5.1.2 Causal Inference

In the causal inference approach, we could consider an individual and define their possible treatment assignment at randomization as $A = 0, 1$ representing placebo or semaglutide respectively. We then define their change from baseline HbA1c values as repeated measures Y_j at visit $j = 1, \dots, J$ with the final change from baseline HbA1c as the outcome of interest Y_J . We also define $D_j = 0, 1$ with value 1 for discontinuation of assigned treatment at visit j and create a vector over the trial visits as $\bar{D} = [D_1, \dots, D_J]$. Similarly, we also define use of rescue therapy as $R_j = 0, 1$ and $\bar{R} = [R_1, \dots, R_J]$. This allows the creation of potential outcomes of $Y_J^{(A, \bar{D}, \bar{R})}$. Using this notation, we could also create a DAG to illustrate our assumptions about the causal relationships between the randomized assignment, treatment discontinuation, use of rescue, and the post-baseline HbA1c values. In some cases, it may also be necessary to consider other post baseline (time-varying) factors. Finally, we can specify the causal estimands that we are interested in via the following mathematical statements about the potential outcomes.

Causal Estimand 1.1

$$\mathbb{E}[Y_J^{(a=1, \bar{D}, \bar{R})}] - \mathbb{E}[Y_J^{(a=0, \bar{D}, \bar{R})}] = \mathbb{E}[Y_J^{(a=1)}] - \mathbb{E}[Y_J^{(a=0)}]$$

As \bar{D} and \bar{R} are not intervened on in this estimand, we can simplify the notation by removing them (last expression). This estimand is then the difference in expected values for a patient assigned to semaglutide vs placebo without regard for discontinuation or use of rescue medication.

Causal Estimand 1.2

$$\mathbb{E}[Y_J^{(a=1, \bar{d}=\mathbf{0}, \bar{r}=\mathbf{0})}] - \mathbb{E}[Y_J^{(a=0, \bar{d}=\mathbf{0}, \bar{r}=\mathbf{0})}]$$

Where $\mathbf{0} = [0, \dots, 0]$ corresponds to a $1 \times J$ vector of zeros. In words, this is targeting the difference in expected outcomes for a patient assigned to semaglutide vs placebo if they received their assigned treatment for the duration of the trial and do not initiate rescue medication. This type of estimand can also be interpreted as a controlled direct effect from causal mediation literature, where the effect is controlled to have no discontinuation from assigned treatment and no initiation of rescue^[41].

5.2 Example 2 - Nasal Polyps

For this example, we only concentrate on the first co-primary endpoint and consider a situation where part of the objective for the SYNAPSE trial was to answer the clinical question:

1. What is the effect of being assigned mepolizumab vs placebo on nasal polyp score at the end of the trial where nasal polyp surgery is considered a poor outcome.

A key issue in the clinical question is how best to represent nasal polyp surgery in terms of the outcome measure nasal polyp score. As requiring surgery is a bad outcome for patients, it is informative about the efficacy provided by either assigned treatment. In this example one option to reflect surgery occurring would be to consider those patients as having the worst possible nasal polyp score.

5.2.1 Estimands Framework

In the estimand framework discontinuing assigned treatment and nasal polyp surgery would be considered IEs. These would then be included as the fifth component of the ICH estimand together with handling strategies that align well with the clinical question of interest above. We provide a potential ICH estimand for the clinical question below:

ICH Estimand 2.1 *Difference in mean nasal polyp Score at 52 weeks for adult patients with chronic rhinosinusitis assigned to 100 mg mepolizumab every four weeks or placebo including any subsequent effects of discontinuation of assigned treatment and considering nasal polyp surgery as a poor outcome reflected as the worst possible nasal polyp score.*

Population	Adult patients with chronic rhinosinusitis
Treatment	100 mg Mepolizumab every four weeks or placebo, including any subsequent discontinuation
Variable	Nasal polyp score at 52 weeks or the worst possible Nasal polyp score if nasal polyp surgery occurs during the 52 weeks.
Summary Measure	The difference in means
IEs (Handling)	Discontinuation of assigned treatment (Treatment Policy) Nasal polyp surgery (Composite)

5.2.2 Causal Inference

In the causal inference approach, we use the same notation as the first example for outcome, assignment, visit and discontinuation of assigned treatment. We also define an individual having nasal polyp surgery by visit j as $S_j = 0, 1$. This allows the creation of two sets of potential outcomes, one for the nasal score $Y_J^{(A, \bar{D})}$, and also another for surgery prior to visit J as $S_J^{(A, \bar{D})}$. We could also create a DAG to communicate our assumptions about the causal relationships between randomized assignment, treatment discontinuation, surgery, and the post-baseline nasal polyp scores. Finally, with notation defined we can describe the following estimands.

Causal Estimand 2.1

For this estimand, there are two components to the effects of interest, first, any direct effect of assignment on the nasal polyp score and second, any indirect but informative effect of assignment on nasal polyp surgery. To reflect this, we construct a new outcome Y_J^* that is a mixture of nasal score potential outcomes for patients when no nasal polyp surgery is performed, and the value chosen to reflect the poor outcome for patients where nasal polyp surgery would be performed. Again, as with causal estimand 1.1, there is no condition applied to \bar{D} and therefore it is not included, and the new outcome becomes:

$$Y_J^{*(A)} = \mathbb{I}(S_J^{(A)} = 0) \cdot Y_J^{(A)} + \mathbb{I}(S_J^{(A)} \neq 0) \cdot \theta$$

Where $\mathbb{I}(\cdot)$ are indicators of the condition contained in the brackets and θ is the value chosen for the individual to represent the poor outcome.

In the SYNAPSE example, we consider a case when an exact value of worst possible nasal polyp score is defined as the value to represent a poor outcome in both arms (θ is set as a score of 4 in each nostril

making a value of 8) and so our resulting estimand at the final timepoint J is:

$$\begin{aligned}\mathbb{E}[Y_J^{*(a=1)}] - \mathbb{E}[Y_J^{*(a=0)}] &= \mathbb{E} \left[\mathbb{I}(S_J^{(a=1)} = 0) \cdot Y_J^{(a=1)} + \mathbb{I}(S_J^{(a=1)} \neq 0) \cdot 8 \right] \\ &\quad - \mathbb{E} \left[\mathbb{I}(S_J^{(a=0)} = 0) \cdot Y_J^{(a=0)} + \mathbb{I}(S_J^{(a=0)} \neq 0) \cdot 8 \right]\end{aligned}$$

Applying the expectations leads to:

$$\begin{aligned}\mathbb{E}[Y_J^{*(a=1)}] - \mathbb{E}[Y_J^{*(a=0)}] &= \left[\mathbb{P}(S_J^{(a=1)} = 0) \cdot \mathbb{E}[Y_J^{(a=1)} | S_J^{(a=1)} = 0] + \mathbb{P}(S_J^{(a=1)} \neq 0) \cdot 8 \right] \\ &\quad - \left[\mathbb{P}(S_J^{(a=0)} = 0) \cdot \mathbb{E}[Y_J^{(a=0)} | S_J^{(a=0)} = 0] + \mathbb{P}(S_J^{(a=0)} \neq 0) \cdot 8 \right]\end{aligned}$$

In words, this estimand is the difference in expected combined nasal polyp score for a patient assigned to mepolizumab vs placebo without regard to treatment discontinuation. The expected combined nasal polyp score under each treatment is the weighted combination of the expected nasal polyp score among those who do not need surgery and the worst possible nasal polyp score, with weights determined by the proportion of patients who need surgery under the assigned treatment. The weighting of the expected values builds in any causal difference in the probability of requiring surgery from being assigned to mepolizumab or placebo.

As discussed in Keene 2019^[42], the penalty value used for estimand 2.1 (defined using ICH or causal inference) is somewhat arbitrary. It also has the consequence of creating an effect that is specific to the θ value(s) used as the penalty term(s). This can make it difficult to compare these types of effects across trials if there is not general agreement on the penalty value.

Importantly, this example also shows that a clinical question of interest may require an estimand that uses different strategies for different IEs.

5.3 Example 3 - COVID-19

For this example, the primary endpoint was occurrence of symptomatic COVID-19 with onset at least 14 days after the second injection and we focus on the key issue that some participants failed to receive two doses of assigned treatment. (Occurrence of symptomatic COVID-19 within 14 days of randomisation represents another relevant issue that we will not discuss here). We consider a situation where the objective was to answer these two different clinical questions:

1. What is the effect of being assigned mRNA-1273 vaccine compared to placebo on COVID-19 prevention if all participants were treated correctly.
2. What is the effect of being assigned mRNA-1273 vaccine compared to placebo on COVID-19 prevention for the participants that would receive treatment correctly when assigned to either treatment.

This example illustrates the subtle difference between interest in an effect on all participants in a hypothetical setting where the IE would not occur for any participants, and an effect specifically for participants where the IE would not have occurred after being assigned either treatment. The former is an effect relating to the whole population, whereas the latter is an effect in a reduced set of participants from the original population.

Although a variety of approaches have been used to derive vaccine efficacy, we concentrate on using a relative risk based measure for vaccine efficacy^[43]

5.3.1 Estimands Framework

In the estimand framework failure to receive two doses of assigned treatment in the correct time window would be considered an IE. This would then be included as the fifth component of an ICH estimand together with a handling strategy that aligns with the clinical questions of interest above. We provide potential ICH estimands for each question below:

ICH Estimand 3.1

The percentage vaccine efficacy (derived from a risk reduction ratio) at 120 days for COVID-19 infection when assigned two 100 mcg mRNA-1273 vaccine injections or placebo in adults with no known history of COVID-19 who are in circumstances that put them at risk of COVID-19 infection assuming that they all received two doses of their assigned treatment in the correct time window.

Population	Adults with no known history of COVID-19 in circumstances that put them at risk of COVID-19 infection
Treatment	Two 100 mcg mRNA-1273 vaccine injections or placebo, in the correct time window
Variable	COVID-19 infection occurring more than 14 days after receiving treatment.
Summary Measure	Vaccine Efficacy at 120 days calculated as percentage risk reduction ratio ($VE\% = 100 \cdot (1 - RR)\%$)
IEs (Handling)	Failure to receive two doses of assigned treatment in the correct time window (Hypothetical)

ICH Estimand 3.2

The percentage vaccine efficacy (derived from a risk reduction ratio) at 120 days for COVID-19 infection when assigned two 100 mcg mRNA-1273 vaccine injections or placebo in adults with no known history of COVID-19 who are in circumstances that put them at risk of COVID-19 infection that would receive two doses of assigned treatment in the correct time window if assigned to either treatment.

Population	Adults with no known history of COVID-19 in circumstances that put them at risk of COVID-19 infection that would receive two doses of treatment assigned to them in the correct time window
Treatment	Two 100 mcg mRNA-1273 vaccine injections or placebo.
Variable	COVID-19 infection occurring more than 14 days after receiving treatment.
Summary Measure	Vaccine Efficacy at 120 days calculated as percentage risk reduction ratio ($VE\% = 100 \cdot (1 - RR)\%$)
IEs (Handling)	Failure to receive two doses of assigned treatment in the correct time window (Principal Stratum)

5.3.2 Causal Inference

In the causal inference approach, we use the same notation as the previous example for outcome, however we note COVID-19 infection is now a binary outcome. We also define a dosing indicator $C = 0, 1$ where $C = 1$ corresponds to receiving the two doses of assigned treatment in the correct time window and this allows the creation of binary potential outcomes $Y^{(A, C)}$ which take the value 1 if the individual had an infection in the time window of 14 days after receiving treatment A . We

could also create a DAG to communicate our assumptions about the casual relationships between the randomized assignment, treatment discontinuation, and COVID-19 infection. With notation defined we can describe the first causal estimand.

Causal Estimand 3.1

$$100 \cdot \left(1 - \frac{\mathbb{E}[Y^{(a=1, c=1)}]}{\mathbb{E}[Y^{(a=0, c=1)}]} \right)$$

In words, this estimand is a percentage reduction of infection calculated using risk ratio of infection if participants were assigned mRNA-1273 vaccine or placebo assuming all participants received both doses of assigned medication in the correct time window. Again, similar to casual estimand 1.2, this type of effect can be interpreted as a controlled direct effect from casual mediation literature, where the effect is controlled to have no dosing irregularities^[41]

Causal Estimand 3.2

For this estimand we extend our dosing notation C to create a potential outcome $C^{(A)}$ for correct dosing under each treatment assignment. This allows us to classify participants into groups based on their potential to receive the two doses correctly after being assigned to each treatment. This creates four potential types (or strata) of participants: Participants that would not receive two doses in the correct time window if assigned to either treatment, participants that would receive two doses in the correct time window if assigned to one treatment, but not the other, and finally, participants that would have received two doses in the correct time window if assigned to either treatment. The last type are the participants forming the principal stratum of interest.

With the notation extended to create the classifications we can write down the second causal estimand.

$$100 \cdot \left(1 - \frac{\mathbb{E}[Y^{(a=1)} \mid C^{(a=0)} = C^{(a=1)} = 1]}{\mathbb{E}[Y^{(a=0)} \mid C^{(a=0)} = C^{(a=1)} = 1]} \right)$$

In words, this effect is the expected percentage reduction in risk of COVID-19 infection for participants assigned mRNA-1273 vaccine or placebo that are members of the stratum that would receive both doses of their assigned treatment regardless of assignment. We note that the definition of this effect is in terms of potential outcomes that cannot both be observed. Therefore, it is only possible to estimate this effect using strong and untestable assumptions. We also note that different opinions may exist on the clinical importance of particular strata. For instance, it could be argued that a stratum of “participants that would receive the mRNA-1273 vaccine correctly” may be considered more relevant. Currently within the literature there are differing opinions of the utility of the principal stratum approach^[44,45]. The strategy also does not exist in all casual ontologies^[46] and therefore using this strategy remains controversial^[47,48].

This example contrasts the difference between effects that apply to the entire population (estimand 3.1) and effects that apply to a reduced number of participants from the original population (estimand 3.2).

5.4 Example 4 - Arthritis

For this example, we concentrate on the comparison of celecoxib and naproxen at the specific doses of 100-200 mg bid and 375-500 mg bid respectively. We focus on the issue of large numbers of patients discontinuing assigned treatment and consider a situation where part of the objective for the trial was to answer the clinical question:

1. What is the difference in cardiac risk when taking celecoxib vs taking naproxen.

The main outcome related to cardiovascular risk was a composite outcome of cardiovascular death, nonfatal myocardial infarction, or nonfatal stroke using the event measure "Antiplatelet Trialists Collaboration (APTC) composite CV endpoint" (APTC-CV). As the clinical focus was the cardiac risk attributable to the treatments in the trial, and a large number of patients discontinued their assigned treatment, it is natural to consider the APTC-CV events occurring while the patients are receiving the treatment. In many cases, this period may include time after the patient has actually stopped taking the medication to reflect the time for clearance of the treatment.

5.4.1 Estimands Framework

In the estimand framework discontinuation of assigned treatment would be considered an IE. This would then be included as the fifth component of an ICH estimand together with a handling strategy that aligns with the clinical questions of interest above. We provide a potential ICH estimand below:

ICH Estimand 4.1

Difference in proportion of patients with the Antiplatelet Trialists Collaboration (APTC) composite CV endpoint (APTC-CV) in adult patients with osteo or rheumatoid arthritis that have concomitant or increased risk of cardiovascular disease that are assigned to celecoxib 100-200 mg bid or naproxen 375-500mg bid while on assigned treatment (defined as randomization up to 30 days after discontinuation from assigned treatment)

Population	Adult patients with osteo or rheumatoid arthritis that have concomitant or increased risk of cardiovascular disease.
Treatment	Celecoxib 100-200 mg bid compared to Naproxen 375-500 mg bid
Variable	Occurrence of Antiplatelet Trialists Collaboration (APTC) composite CV endpoint while on assigned treatment (defined as randomization up to 30 days after discontinuation).
Summary Measure	Proportion of patients experiencing APTC-CV while on randomized treatment.
IEs (Handling)	Discontinuation of assigned treatment (While on Treatment).

5.4.2 Causal Inference

In the causal inference approach, we use the same notation as the previous examples for treatment assignment. We define the outcome of APTC-CV occurring before time t as $Y(t)$, the time of discontinuation from assigned treatment as T_D , and the (fixed) length of the trial as t_{max} . To include the clearance of treatments, T_D is defined as 30 days after the actual time a participant discontinues assigned treatment. With notation defined we can describe the causal estimand below.

Causal Estimand 4.1

For this estimand there are two components to the effect of interest, any direct effect of assignment on the risk of patients experiencing APTC-CV events, and any indirect effect on the time spent at risk of APTC-CV events due to discontinuation of assigned treatment. To reflect this, we create a binary potential outcome for experiencing an APTC-CV event before time t as $Y^{(A)}(t)$ and the time to discontinuation of randomized treatment or end of follow up time as $T^{(A)} = \min(T_D^{(A)}, t_{max})$. The estimand is then defined as:

$$\mathbb{E}[Y_J^{(a=1)}(T^{(a=1)})] - \mathbb{E}[Y_J^{(a=0)}(T^{(a=0)})]$$

In words this estimand is the difference in expected proportion experiencing an APTC-CV event while on assigned treatment for patients assigned to Celecoxib compared to Naproxen.

This estimand has clear interpretation when the time to discontinuation of randomized treatment is similar in the two arms (as was the case in the example trial). However, if the discontinuation rates were to differ according to treatment assignment, the interpretation becomes more complex due to the differential time ‘at-risk’ between the arms. In these scenarios a clearer estimand could be formed where the number of APTC-CV events is considered relative to the assigned treatment exposure time^[49,50].

6 Discussion

6.1 Comparing the ICH and Causal Paradigms

The examples above demonstrate how both the estimand framework and causal inference methodology can be used to define estimands. Both approaches allow clinical teams flexibility in defining the specific effect of interest and facilitate the discussion about which overall effect is being considered in light of various complicating factors. The ICH estimand is explicit with five components that must be defined, whereas the causal estimand is the result of defining the ‘treatment’ variables and potential outcomes to be compared and the assumed causal structures impacting the outcome (often via a DAG).

For causal estimands, factors such as treatment assignment, stopping treatment, and using additional or rescue medication are specified in the algebra which contains a considerable amount of information in a very subtle way. It is precise, but the concepts it relies on, and subtleties of specification may be harder to comprehend. This mathematical expression, together with a well-defined population and detail on the outcome of interest, provides the same overall information as an ICH estimand. Further, the use of DAGs allows the assumed causal structure of the outcome, assignment, and other factors to be communicated to guide the estimation of the effect targeted.

In contrast, the ICH estimand is not mathematical and is expressed either in a formal table or a single dense paragraph of text. This is likely to be more accessible to clinical trialists without formal causal and mathematical training. However, because the text is not as precise as the causal algebra, it is important to define the ICH estimand using detailed statements.

As demonstrated in our examples, the concept of an intercurrent event (IE) and strategies to address them borrow ideas and concepts from causal inference. It is reasonably intuitive to consider the concept

of an IE to be the same as post-baseline mediator or time varying treatment developed as part of the causal inference literature.

In hindsight, given the causal inference and biostatistics literature overlap heavily, and the term estimand has existed for many years, it might have been helpful for ICH E9 (R1) to have discussed linkages to causal inference terms and concepts explicitly. We think the creation of the ICH estimands framework was needed to provide clarity on what is estimated in clinical trials, and all five components are required to gain the understanding of what is being estimated, but the lack of discussion linking it to established causal concepts has created some confusion.

Finally, the examples demonstrate that what is called an estimand in the two paradigms may only partially overlap but are essentially both targeting the same thing: a population-based summary of a causal effect on an outcome in a particular population.

6.2 General Comments on Causal Inference

In this work we included a description of the specific elements of causal inference required to compare causal and ICH estimands in the context of RCTs. The field of causal inference is much broader than the areas introduced in this paper and developed mainly in the context of analysis of observational data, where careful consideration of confounding (and other complicating factors) is required to enable sound analytical conclusions. The complex nature of observational data inspired the development of methods like the potential outcomes framework and DAGs, which in turn led directly to the development of other areas such as target trial emulation^[51], Single World Intervention Graphs (SWiGs)^[52], and causal discovery^[53].

The causal inference community have also improved the standard of clinical trial methodology; examples include highlighting the need to distinguish between conditional and marginal treatment effects^[54] and warning about interpretational subtleties that can arise when using hazard ratios^[55–57].

When applying causal inference to clinical trials in drug development, generally only two treatment effects of interest have commonly been identified: (i) the effect of assignment to the intervention and (ii) the effect of adhering to the intervention as specified in the trial protocol^[58–60]. The latter is often described as the per protocol effect, specifically “...*the causal effect of treatment that would have been observed if all individuals had adhered to their assigned treatment as specified in the protocol of the experiment.*”^[6,59]. The term “per-protocol effect” may be confusing to non-statisticians in drug development since it is often interpreted incorrectly as an analysis of a per-protocol population^[61]. Similarly, the term “adherence” is sometimes ill-defined and can refer to continuing to take a medication or to some formal measure of compliance. It is also unclear in the causal inference literature how events such as use of rescue medications or surgery are accounted for when discussing adherence, particularly in trials where these additional interventions are indicated in the protocol.

While “assignment” or “adherence” effects may answer important scientific questions, there are other research questions that a multi-stakeholder trial would regularly aim to answer and contrary to this perceived dichotomy, causal inference allows the definition of effects that are not purely from assignment or perfect protocol compliance. Therefore, suggestions that the assignment and adherence effects are the only important questions is unhelpful^[61].

6.3 General Comments on on ICH E9 (R1)

Although finalized in 2019, ICH E9(R1) still contains some ambiguity on certain points. As a result, further work has been published suggesting ICH estimands should have certain additional properties^[62]. We comment on some aspects of the addendum in general below and would welcome further discussion with regulators.

In general, the label of “treatment policy” is not particularly accurate when this strategy is used for handling discontinuation of assigned treatments. In many diseases, when the randomized treatment fails, the actual treatments participants receive would be at the discretion of the physician and would typically contain new treatments or no treatment at all. Any follow-on treatments could also vary substantially depending on factors such as geographical region or the trial time. In these cases, there is not a well-defined “treatment policy” and the strategy itself would be better labelled as the assessing the effect of assignment.

A potentially confusing aspect of the addendum relates to the principal stratum strategy. ICH E9(R1) emphasises that this strategy did not just correspond to an analysis of a post baseline subgroup, specifying: *“It is important to distinguish “principal stratification” (see Glossary), which is based on potential intercurrent events (for example, subjects who would discontinue therapy if assigned to the test product), from sub setting based on actual intercurrent events (subjects who discontinue therapy on their assigned treatment).”* However, when an intercurrent event is not affected by treatment assignment, Kahan et al.^[63,64] state that valid estimates for a principal stratum can be obtained by comparing randomised groups that exclude those who experienced the intercurrent event (i.e. comparing post baseline subgroups). In many cases, it may not be plausible to assume the randomized treatment has no effect on intercurrent events and this analysis approach would then not be justified.

There are also contrasting opinions on the clinical relevance of the principal stratum strategy as individual membership cannot be identified at baseline^[44,45].

Although the hypothetical strategy can be viewed as the effect if the IE had not occurred^[3], non-statisticians sometimes find the term “hypothetical” problematic. This is because the term suggests something is being estimated that is different from factual evidence and, for this reason, the strategy can be viewed as irrelevant to use of the medicine in clinical practice^[65]. However, applying this strategy to IEs in many cases, attempts to assess the causal effect of the medicine itself and this effect may be relevant to patients and prescribers. An effect described in the product description (label) for a medication based on a treatment policy approach will reflect a mixture of results from participants taking a range of possible treatments. Clear identification of the causal effect of the treatment itself would be helpful for patients that are considering whether to take the medication.

A potential approach to addressing this is the attributable estimand described by Darken et al.^[65]. In this approach, events that are considered to be adversely related to randomised treatment (e.g. discontinuation of treatment due to adverse events or lack of efficacy) are considered attributable and handled with a composite strategy (hence become part of the endpoint of interest), while a hypothetical strategy is used for intercurrent events not considered to be related to randomised treatment (e.g. treatment discontinuation due to administrative reasons such as loss to follow-up). Estimates of treatment effect from an attributable estimand may be of more use to patients and prescribers than the common practice of providing estimates based on a treatment policy strategy.

7 Conclusion

We have discussed estimands as defined in ICH E9 (R1) as part of the “estimand framework” and causal inference field. We provided a brief commentary on where both concepts of an estimand originated and used real trial examples with clinical questions of interest to illustrate how both ICH and causal estimands could be formulated. We then compared the similarities and differences between the approaches.

We conclude both can be used to quantify a population-based summary of a target effect on an outcome for a given population. Both paradigms rely on assumptions that underpin causal inference, and both have their use in the field of clinical trials. The ICH estimand is a description-based definition which is likely to be more accessible to non-mathematicians. The causal estimand is more precise due to its mathematical nature but may be harder to grasp for people without the causal training to mathematically define the target of inference and communicate their assumptions through causal graph diagrams.

Although we have presented both interpretations of estimands side by side for understanding and comparison, we would encourage practitioners to consider them as complementary and not competing frameworks. Both approaches have attractive features and combining them allows the trialist more options for communicating what is being estimated in clinical trials. We encourage people familiar with the estimands framework to use causal inference and vice versa.

8 Acknowledgement

The authors would like to thank Matt Psioda and Ben Hartley for discussions on potential outcomes and composite IE handling strategies.

9 Conflicts of Interest

TD owns shares in GlaxoSmithKline. JWB and his employers have received fees for statistical consultancy from AstraZeneca, Bayer, Novartis, and Roche. DW owns shares in AstraZeneca. ONK owns shares in GlaxoSmithKline and has received fees for statistical consultancy from AstraZeneca and Areteia.

References

- [1] ICH. E9(R1) (2019) Addendum on Estimands and Sensitivity Analysis in Clinical Trials. <https://www.ich.org/page/efficacy-guidelines> 2019. [Accessed 06-12-2024].
- [2] Mallinckrodt C H, Bell J, Liu G, et al. Aligning estimators with estimands in clinical trials: Putting the ICH E9(R1) guidelines into practice *Ther. Innov. Regul. Sci.*. 2020;54:353–364.
- [3] Fletcher C, Hefting N, Wright M, et al. Marking 2-years of new thinking in Clinical Trials: The estimand journey *Ther. Innov. Regul. Sci.*. 2022;56:637–650.

- [4] Cro Suzie, Kahan Brennan C, Rehal Sunita, et al. Evaluating how clear the questions being investigated in randomised trials are: systematic review of estimands *BMJ*. 2022;378:e070146.
- [5] ICH. E9(R1) Training Material. https://database.ich.org/sites/default/files/E9%28R1%29%20Training%20Material%20-%20PDF_0.pdf 2020. [Accessed 06-12-2024].
- [6] Hernan Miguel A, Robins James M. *Causal inference: What if*. Boca Raton, FL: CRC Press 2024.
- [7] Estimands and Causal Inference https://psiweb.org/docs/default-source/default-document-library/daniel-scharfstein-slides.pdf?sfvrsn=28c8dedb_0. 2017. [Accessed 13Nov2024].
- [8] Scharfstein Daniel. A constructive critique of the draft ICH E9 Addendum *Clinical Trials*. 2019;16:375-380.
- [9] Lipkovich Ilya, Ratitch Bohdana, Mallinckrodt Craig H.. Causal Inference and Estimands in Clinical Trials *Statistics in Biopharmaceutical Research*. 2020;12:54-67.
- [10] Han Shasha, Zhou Xiao-Hua. Defining estimands in clinical trials: A unified procedure *Statistics in Medicine*. 2023;42:1869-1887.
- [11] Van Lancker Kelly, Vo Tat-Thang, Akacha Mouna. Estimands in heath technology assessment: a causal inference perspective *Statistics in Medicine*. 2022;41:5577-5585.
- [12] Bowden Jack, Bornkamp Björn, Glimm Ekkehard, Bretz Frank. Connecting Instrumental Variable methods for causal inference to the Estimand Framework *Statistics in Medicine*. 2021;40:5605-5627.
- [13] Aroda Vanita R, Rosenstock Julio, Terauchi Yasuo, et al. PIONEER 1: Randomized clinical trial of the efficacy and safety of oral semaglutide monotherapy in comparison with placebo in patients with type 2 diabetes *Diabetes Care*. 2019;42:1724-1732.
- [14] Han Joseph K, Bachert Claus, Fokkens Wytse, et al. Mepolizumab for chronic rhinosinusitis with nasal polyps (SYNAPSE): a randomised, double-blind, placebo-controlled, phase 3 trial *Lancet Respir. Med.*. 2021;9:1141-1153.
- [15] Baden Lindsey R, El Sahly Hana M, Essink Brandon, et al. Efficacy and safety of the mRNA-1273 SARS-CoV-2 vaccine *N. Engl. J. Med.*. 2021;384:403-416.
- [16] Nissen Steven E., Yeomans Neville D., Solomon Daniel H., et al. Cardiovascular Safety of Celecoxib, Naproxen, or Ibuprofen for Arthritis *New England Journal of Medicine*. 2016;375:2519-2529.
- [17] National Research Council (US) Panel on Handling Missing Data in Clinical Trials: The Prevention and Treatment of Missing Data in Clinical Trials. <https://www.ncbi.nlm.nih.gov/books/NBK209902/> 2010. [Accessed 06-12-2024].
- [18] Ranganathan Priya, Pramesh C S, Aggarwal Rakesh. Common pitfalls in statistical analysis: Intention-to-treat versus per-protocol analysis *Perspect. Clin. Res.*. 2016;7:144-146.
- [19] Hernán Miguel A, Hernández-Díaz Sonia. Beyond the intention-to-treat in comparative effectiveness research *Clin. Trials*. 2012;9:48-55.
- [20] Lachin J M. Statistical considerations in the intent-to-treat principle *Control. Clin. Trials*. 2000;21:167-189.
- [21] Ellenberg Jonas H. Intent-to-treat analysis versus as-treated analysis *Drug Inf. J.*. 1996;30:535-544.
- [22] Lynggaard Helle, Bell James, Lösch Christian, et al. Principles and recommendations for incorporating estimands into clinical study protocol templates *Trials*. 2022;23:685.

- [23] Roger James H, Bratton Daniel J, Mayer Bhabita, Abellan Juan J, Keene Oliver N. Treatment policy estimands for recurrent event data using data collected after cessation of randomised treatment *Pharm. Stat.*. 2019;18:85–95.
- [24] Polverejan Elena, Dragalin Vladimir. Aligning treatment policy estimands and estimators—A simulation study in Alzheimer’s disease *Stat. Biopharm. Res.*. 2020;12:142–154.
- [25] Guizzaro Lorenzo, Pétavy Frank, Ristl Robin, Gallo Ciro. The use of a variable representing compliance improves accuracy of estimation of the effect of treatment allocation regardless of discontinuation in trials with incomplete follow-up *Stat. Biopharm. Res.*. 2021;13:119–127.
- [26] Hartley Benjamin, Drury Thomas, Lettis Sally, Mayer Bhabita, Keene Oliver N, Abellan Juan J. Estimation of a treatment policy estimand for time to event data using data collected post discontinuation of randomised treatment *Pharm. Stat.*. 2022;21:612–624.
- [27] Wolbers Marcel, Noci Alessandro, Delmar Paul, Gower-Page Craig, Yiu Sean, Bartlett Jonathan W. Standard and reference-based conditional mean imputation *Pharm. Stat.*. 2022;21:1246–1257.
- [28] Wang Yun, Tu Wenda, Kim Yoonhee, et al. Statistical methods for handling missing data to align with treatment policy strategy *Pharm. Stat.*. 2023;22:650–670.
- [29] Noci Alessandro, Wolbers Marcel, Abt Markus, et al. A comparison of estimand and estimation strategies for clinical trials in early Parkinson’s disease *Stat. Biopharm. Res.*. 2022:1–22.
- [30] Drury Thomas, Abellan Juan J, Best Nicky, White Ian R. Estimation of treatment policy estimands for continuous outcomes using off-treatment sequential multiple imputation *Pharm. Stat.*. 2024;23:1144–1155.
- [31] Bell James, Drury Thomas, Mütze Tobias, et al. Estimation Methods for Estimands Using the Treatment Policy Strategy; a Simulation Study Based on the PIONEER 1 Trial *Pharmaceutical Statistics*. 2025;24:e2472. e2472 PST-24-0023.R1.
- [32] Olarte Parra Camila, Daniel Rhian M, Bartlett Jonathan W. Hypothetical estimands in clinical trials: A unification of causal inference and missing data methods *Stat. Biopharm. Res.*. 2023;15:421–432.
- [33] Rubin Donald B. Causal Inference Using Potential Outcomes: Design, Modeling, Decisions *Journal of the American Statistical Association*. 2005;100:322–331.
- [34] Pearl Judea. An Introduction to Causal Inference *The International Journal of Biostatistics*. 2010;6.
- [35] Dawid A. P.. Causal Inference without Counterfactuals *Journal of the American Statistical Association*. 2000;95:407–424.
- [36] Pearl Judea. Interview with Judea Pearl *Observational Studies*. 2022;8:23–36.
- [37] Robins Jamie. Interview with Jamie Robins *Observational Studies*. 2022;8:37–76.
- [38] Rubin Don. Interview with Don Rubin *Observational Studies*. 2022;8:77–94.
- [39] Tsiatis Anastasios A., Davidian Marie, Holloway Shannon T., Laber Eric B.. *Dynamic Treatment Regimes: Statistical Methods for Precision Medicine*. Chapman and Hall/CRC 2019.
- [40] Greenland Sander, Pearl Judea, Robins James M.. Causal Diagrams for Epidemiologic Research *Epidemiology*. 1999;10:37–48.
- [41] Robins J M, Greenland S. Identifiability and exchangeability for direct and indirect effects *Epidemiology*. 1992;3:143–155.
- [42] Keene Oliver N.. Strategies for composite estimands in confirmatory clinical trials: Examples from trials in nasal polyps and steroid reduction *Pharmaceutical Statistics*. 2019;18:78–84.

- [43] Orenstein W A, Bernier R H, Dondero T J, et al. Field evaluation of vaccine efficacy *Bull. World Health Organ.*. 1985;63:1055–1068.
- [44] Stensrud Mats J., Dukes Oliver. Translating questions to estimands in randomized clinical trials with intercurrent events *Statistics in Medicine*. 2022;41:3211–3228.
- [45] Bornkamp Björn, Rufibach Kaspar, Lin Jianchang, et al. Principal stratum strategy: Potential role in drug development *Pharmaceutical Statistics*. 2021;20:737–751.
- [46] Sarvet Aaron L, Stensrud Mats J. Without commitment to an ontology, there could be no causal inference *Epidemiology*. 2022;33:372–378.
- [47] Dawid Philip, Didelez Vanessa. “Imagine a can opener”—the magic of principal stratum analysis *Int. J. Biostat.*. 2012;8:19.
- [48] Vansteelandt Stijn, Van Lancker Kelly. Chasing shadows: How implausible assumptions skew our understanding of causal estimands *Stat. Biopharm. Res.*. 2024;1–7.
- [49] Stegherr Regina, Schmoor Claudia, Beyersmann Jan, et al. Survival analysis for Adverse events with VarYing follow-up times (SAVVY)-estimation of adverse event risks *Trials*. 2021;22:420.
- [50] Siddiqui Ohidul. Statistical methods to analyze adverse events data of randomized clinical trials *J. Biopharm. Stat.*. 2009;19:889–899.
- [51] Hernán Miguel A, Wang Wei, Leaf David E. Target trial emulation: A framework for causal inference from observational data *JAMA*. 2022;328:2446–2447.
- [52] Ocampo Alex, Bather Jemar R. Single-world intervention graphs for defining, identifying, and communicating estimands in clinical trials *Stat. Med.*. 2023;42:3892–3902.
- [53] Huber Martin. An introduction to causal discovery *Schweiz. Z. Volkswirtschaft. Stat.*. 2024;160.
- [54] Van Lancker Kelly, Bretz Frank, Dukes Oliver. Covariate adjustment in randomized controlled trials: General concepts and practical considerations *Clin. Trials*. 2024;21:399–411.
- [55] Hernán Miguel A. The hazards of hazard ratios *Epidemiology*. 2010;21:13–15.
- [56] Kalbfleisch J D, Prentice R L. Estimation of the Average Hazard Ratio *Biometrika*. 1981;68:105–112.
- [57] Aalen Odd O, Cook Richard J, Røysland Kjetil. Does Cox analysis of a randomized survival study yield a causal treatment effect? *Lifetime Data Anal.*. 2015;21:579–593.
- [58] Sterne Jonathan A C, Savović Jelena, Page Matthew J, et al. RoB 2: a revised tool for assessing risk of bias in randomised trials *BMJ*. 2019;l4898.
- [59] Hernán Miguel A., Robins James M.. Per-Protocol Analyses of Pragmatic Trials *New England Journal of Medicine*. 2017;377:1391–1398.
- [60] Higgins Julian P.T., Thomas James, Chandler Jacqueline, et al. , eds. *Cochrane Handbook for Systematic Reviews of Interventions*. Wiley 2019.
- [61] Keene Oliver. Adherence, per-protocol effects, and the estimands framework *Pharmaceutical Statistics*. 2023;22:1141–1144.
- [62] Mütze Tobias, Bell James, Englert Stefan, et al. Principles for Defining Estimands in Clinical Trials—A Proposal *Pharmaceutical Statistics*. 2024.
- [63] Kahan Brennan C, White Ian R, Edwards Mark, Harhay Michael O. Using modified intention-to-treat as a principal stratum estimator for failure to initiate treatment *Clin. Trials*. 2023;20:269–275.
- [64] Kahan Brennan C, Hindley Joanna, Edwards Mark, Cro Suzie, Morris Tim P. The estimands framework: a primer on the ICH E9(R1) addendum *BMJ*. 2024;384:e076316.
- [65] Darken Patrick, Nyberg Jack, Ballal Shaila, Wright David. The attributable estimand: A new approach to account for intercurrent events *Pharmaceutical Statistics*. 2020;19:626–635.