

UC Berkeley

UC Berkeley Previously Published Works

Title

Defining and estimating effects in cluster randomized trials: A methods comparison

Permalink

<https://escholarship.org/uc/item/7rv3x83w>

Journal

Statistics in Medicine, 42(19)

ISSN

0277-6715

Authors

Benitez, Alejandra
Petersen, Maya L
van der Laan, Mark J
et al.

Publication Date

2023-08-30

DOI

10.1002/sim.9813

Peer reviewed



Published in final edited form as:

Stat Med. 2023 August 30; 42(19): 3443–3466. doi:10.1002/sim.9813.

Defining and estimating effects in cluster randomized trials: A methods comparison

Alejandra Benitez¹, Maya L. Petersen², Mark J. van der Laan², Nicole Santos³, Elizabeth Butrick³, Dilys Walker³, Rakesh Ghosh³, Phelgona Otieno⁴, Peter Waiswa⁵, Laura B. Balzer²

¹Genentech Inc., South San Francisco, California

²School of Public Health, Biostatistics, University of California Berkeley, Berkeley, California

³Institute for Global Health Sciences, University of California San Francisco, San Francisco, California

⁴Center for Clinical Research, Kenya Medical Research Institute, Nairobi, Kenya

⁵Centre of Excellence for Maternal, Newborn and Child Health, Makerere University College of Health Sciences, Kampala, Uganda

Abstract

Across research disciplines, cluster randomized trials (CRTs) are commonly implemented to evaluate interventions delivered to groups of participants, such as communities and clinics. Despite advances in the design and analysis of CRTs, several challenges remain. First, there are many possible ways to specify the causal effect of interest (eg, at the individual-level or at the cluster-level). Second, the theoretical and practical performance of common methods for CRT analysis remain poorly understood. Here, we present a general framework to formally define an array of causal effects in terms of summary measures of counterfactual outcomes. Next, we provide a comprehensive overview of CRT estimators, including the *t*-test, generalized estimating equations (GEE), augmented-GEE, and targeted maximum likelihood estimation (TMLE). Using finite sample simulations, we illustrate the practical performance of these estimators for different causal effects and when, as commonly occurs, there are limited numbers of clusters of different sizes. Finally, our application to data from the Preterm Birth Initiative (PTBi) study demonstrates the real-world impact of varying cluster sizes and targeting effects at the cluster-level or at the individual-level. Specifically, the relative effect of the PTBi intervention was 0.81 at the cluster-level, corresponding to a 19% reduction in outcome incidence, and was 0.66 at the individual-level, corresponding to a 34% reduction in outcome risk. Given its flexibility to estimate a variety of user-specified effects and ability to adaptively adjust for covariates for precision gains while maintaining Type-I error control, we conclude TMLE is a promising tool for CRT analysis.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Correspondence Alejandra Benitez, Genentech Inc., South San Francisco, CA, USA. alelizbenitez@gmail.com.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

Keywords

cluster randomized trials; clustered data; data-adaptive adjustment; group randomized trials; Hierarchical data; targeted maximum likelihood estimation

1 | INTRODUCTION

Cluster randomized trials (CRTs) provide an opportunity to assess the population-level effects of interventions that are randomized to groups of individuals, such as communities, clinics, or schools. These groups are commonly called clusters. The choice to randomize clusters, instead of individuals, is often driven by the type of intervention as well as practical considerations.¹ For example, interventions to improve medical practices are often randomized at the hospital or clinic-level to reduce logistical burden and to minimize potential contamination between arms if individual patients were instead randomized. The design and conduct of CRTs has improved considerably,²⁻⁵ and results from CRTs have been widely published in public health, education, policy, and economics literature.⁶ However, a recent review found that only 50% of CRTs were analyzed with appropriate methods.⁷

Due to the hierarchical nature of the data and the correlation of participant outcomes within clusters, the analysis of CRTs is fundamentally more complicated than for individually randomized trials.^{1,7} To start, there are many ways to define the causal effect of interest in CRTs. We may, for example, be interested in the effect for the sample enrolled in the CRT or for a wider target population. Furthermore, we may be interested in the effect at the individual-level or cluster-level. As detailed below, the individual- and cluster-level effects can diverge markedly under “informative cluster size,”^{8,9} occurring when cluster size modifies the intervention effect. Finally, various summary measures of arm-specific outcomes (eg, weighted means) and their contrasts (eg, difference or ratio) may be of interest. Altogether, the causal effect should be determined by the study’s goal and primary research question.¹⁰⁻¹³ This is in line with the International Council for Harmonization (ICH)-E9(R1) guidance for trial protocols to explicitly state the “estimand,” including the target population, comparison conditions, endpoint, and summary measure.¹⁴ Adhering to this guidance ensures the statistical estimator follows from the target estimand, which follows from the trial’s objective. Neglecting this guidance risks letting the statistical approach determine the effect estimated and thus determine which research question is answered.

For the analysis of CRTs, statistical estimation and inference must respect that the cluster is the independent unit. Ignoring the dependence of observations within a cluster can lead to power calculations based on the wrong sample size and underestimates of the SE. Analyses ignoring clustering may also inappropriately attribute the impact of a cluster-level covariate to the intervention and, together with variance underestimation, result in inflated Type-I error rates. Many analytic approaches are available to account for the dependence within a cluster; examples include conducting a cluster-level analysis or applying a correction factor.^{1-4,7} Once we have committed to an analytic approach that aligns with our research question and addresses clustering, the adjustment of baseline covariates is often considered

to improve precision and, thereby, statistical power. In an individually randomized trial setting, Benkeser et al recently demonstrated an 18% savings in sample size that could be achieved by including baseline covariates in the analysis.¹⁵ Likewise, using real data from a CRT, Balzer et al recently showed the adjusted analysis was five times more efficient than an unadjusted analysis for the same parameter.¹⁶ While our focus is on the potential of covariate adjustment to improve precision, we note that in other settings covariate adjustment may be essential to reducing bias due to missingness, selection, or restricted randomization.¹⁶⁻²⁵

Fortunately, many methods are available to incorporate covariates for improved efficiency when estimating intervention effects in CRTs. Examples include well-established methods, such as generalized estimating equations (GEE) and covariate-adjusted residuals estimator (CARE), as well as more recent developments, such as targeted maximum likelihood estimation (TMLE) and Augmented-GEE (Aug-GEE).^{1,11,26-30} While these methods differ in their exact implementation, each aims to improve the CRT's statistical power by controlling for individual- or cluster-level covariates when fitting the "outcome regression": the conditional expectation of the outcome given the randomized intervention and the adjustment variables. As detailed below, these algorithms naturally estimate distinct causal effects, again highlighting the importance of pre-specifying the target estimand and then choosing the optimal estimator of that effect. Previous literature has used simulation studies to compare the attained power and Type-I error rate of various approaches to CRT analysis (eg, References 31 and 32). However, to the best of our knowledge, these comparisons have largely excluded the more recent approaches of Aug-GEE and TMLE.

This paper aims to provide a general framework for defining causal effects in CRTs and to assess the comparative performance of analytic methods for estimation of those effects. Building on prior work in CRT analysis (eg, References 1,5,21,24,25,32-34), our key contributions are as follows. First, we demonstrate the utility of a nonparametric structural causal model, accounting for clustering, to derive counterfactuals and define a variety of causal effects of potential interest. Using this causal model framework, we examine the impact of varying cluster sizes when defining the target causal effect and discuss identification of each causal effect as a function of the observed data distribution (ie, a statistical estimand). Next, we review recently developed CRT estimators along with well-established ones, emphasizing their natural target of inference. Specifically, we describe each algorithm's ability to estimate marginal effects on the additive or relative scale and at the individual- or cluster-level, while also adjusting for covariates to improve precision. To the best of our knowledge, this is the first paper to describe several implementations of TMLE for the analysis of CRTs; these TMLEs allow for estimation of individual-level effects with cluster-level data and, conversely, for estimation of cluster-level effects with individual-level data. Also to the best of our knowledge, this is the first paper using finite sample simulations to provide a head-to-head comparison of these TMLEs, overall and under informative cluster size. Importantly, this work is motivated by a real data application with 20 clusters of widely variable size.

For an alternative and complementary presentation using potential outcomes and a "design-based" perspective, we refer the reader to Su and Ding,³² who focus on effects defined

for the sample and on the absolute scale (ie, difference in mean outcomes) as well as estimation using weighted least squares regression. Here, we define and estimate effects on both the absolute and relative scale and for both the sample population and the larger target population. Additionally, our work is applicable regardless of the outcome-type (ie, binary, count, or continuous) and to trials with more complex designs (eg, pair-matched). While prior work has focused on asymptotic theory and provided simulation results with a fairly large number of clusters, here, we focus on the practical application of these estimators and evaluate their performance in simulations with limited numbers of clusters.

As motivating example, we consider the Preterm Birth Initiative (PTBi) study, a maternal-infant CRT which took place in 20 health facilities across Kenya and Uganda ([ClinicalTrials.gov: NCT03112018](https://clinicaltrials.gov/ct2/show/study/NCT03112018)). The trial assessed whether a facility-based intervention, designed to improve the uptake of evidence-based practices, was effective in reducing 28-day mortality among preterm infants. An important feature of PTBi was the widely varying cluster sizes; specifically, the median number of mother-infant dyads within a facility was 236 and ranged from 29 to 447. Following PTBi, we focus on a setting where individual participants are grouped into clusters (eg, patients in a health facility). However, our discussion and results are equally applicable to other hierarchical data structures that may be observed in a CRT. Examples include households within a community, classrooms within a school, and hospitals within a district.

The remainder of the paper is organized as follows. In Section 2, we use a nonparametric structural equation model to provide a broad approach to formally defining causal effects in CRTs. We discuss how the primary endpoint may be defined at the individual-level or at the cluster-level, often as an aggregate of individual-level outcomes. We highlight how these endpoints can differ in terms of interpretation and magnitude, especially under the setting of informative cluster size. We additionally address the distinction between effects defined for the study sample vs a target population. In Section 3, we discuss several CRT estimation methods. For each method, we describe its target of inference and its use of baseline covariates to improve statistical precision. In Section 4, we provide two simulation studies to evaluate the finite sample performance of common CRT estimators and to demonstrate the distinction between different causal effects. In Section 5, we apply these methods to estimate intervention effects for the PTBi study and highlight the real-world impact of varying cluster sizes. We conclude in Section 6 with a brief discussion.

2 | DEFINING CAUSAL EFFECTS IN CRTS

We begin by formalizing the notation that will be used throughout. We index the cluster, such as a community or hospital, by $j = 1, \dots, J$. The study clusters could be randomly sampled from a larger target population of clusters or, due to practical constraints, selected for convenience from a set of candidate clusters. In either case, there is a real or theoretical target population of clusters to which we *may* want to make inferences. Below, we explicitly discuss the definition and estimation of the population versus sample effects. In all settings, the selection of clusters should be reflected in the CONSORT diagram. Each cluster is comprised of a finite set of individuals (ie, participants), which are indexed $i = 1, \dots, N_j$. The cluster size, denoted N_j , could be constant across clusters, but more often varies by cluster.

The N_j study participants could be randomly sampled within a cluster or be a census of all persons in a cluster. Discussion of more complex settings with systematic sampling or other forces of selection bias is important but beyond the scope of this paper.^{16,20-23} Here, the study participants are representative of or comprise the cluster. Thus, for a given cluster, N_j is fixed and finite.

In each cluster, $\mathbf{W}_j = (W_{1j}, \dots, W_{N_jj})$ denotes the set of baseline covariates for the study participants. Additional baseline covariates related to the cluster are denoted E_j ; these could be summaries of individual-level characteristics or may have no individual-level counterpart.²⁸ For ease of notation, we include the cluster size N_j , a random variable, in E_j . These cluster-level and individual-level covariates are assumed not to be impacted by the intervention, which is randomized to clusters in a CRT. In the case of a pair-matched trial, which may increase study power,^{1,33,35} randomization occurs within pairs of clusters matched on characteristics expected to be predictive of the primary endpoint. Throughout, we will use A_j as an indicator that cluster j was randomized to the intervention. The outcome of interest is $\mathbf{Y}_j = (Y_{1j}, \dots, Y_{N_jj})$, which for simplicity we assume is measured for all participants in cluster j . Extensions to handling missing data time-dependent covariates, censoring, and selection bias are important, but beyond the scope of this paper.^{16-18,20-23}

2.1 | Hierarchical structural causal model

We now use Pearl's nonparametric structural equation model³⁶ to represent the hierarchical data generating process of a CRT. As detailed in Balzer et al,²⁸ the following model encodes independence between clusters, but makes no restrictions on the dependence of participants within a cluster.

$$\begin{aligned} E &= f_E(U_E) \\ \mathbf{W} &= f_W(E, U_W) \\ A &= f_A(U_A) \\ \mathbf{Y} &= f_Y(E, \mathbf{W}, A, U_Y). \end{aligned} \tag{1}$$

Here, (f_E, f_W, f_A, f_Y) denote the set of functions that determine the values of the measured random variables: the cluster-level covariates E , the individual-level covariate matrix \mathbf{W} , the cluster-level intervention assignment A , and the vector of individual-level outcomes \mathbf{Y} . These functions also account for unmeasured factors (U_E, U_W, U_A, U_Y) . By design in a CRT, the unmeasured factors determining the intervention assignment U_A are independent of others. These functions are also nonparametric with the exception of f_A , the function to generate the cluster-level intervention. For example, in a two-armed trial with equal allocation probability, we have $f_A = \mathbb{1}(U_A < 0.5)$ with $U_A \sim Unif(0, 1)$. In contrast, the function to generate the outcome vector f_Y is left unspecified and can take any form of the "parent" variables (E, \mathbf{W}, A) and unmeasured factors U_Y . Importantly, the cluster size N , included in E , can influence the outcome vector \mathbf{Y} in complex ways. For example, very large and very small clusters may have poorer outcomes. Additionally the intervention effect maybe attenuated or, alternatively, enhanced in the largest and smallest clusters. Beyond the cluster-

level covariates and intervention (E, A), there are additional unmeasured and measured factors influencing and inducing dependence in the outcomes \mathbf{Y} in a cluster. For example, the joint error term U_Y induces correlation among participants' outcomes within a cluster. Additionally, within cluster j , an individual's outcome Y_{ij} may depend on the covariates of others in the same cluster \mathbf{W}_j . In other settings, the dependence between participants in a cluster might be more restricted; for further details, we refer to reader to Balzer et al.²⁸

We now consider the data generating process for the PTBi study, where the cluster corresponds to a health facility and the "individual" to a mother-infant dyad. For cluster j , we measure facility-level baseline characteristics E_j , including the average monthly delivery volume, facility preparedness assessment score, staff to delivery ratio, and community-type (ie, urban vs rural). The facility is then randomly assigned to intervention ($A_j = 1$) or control ($A_j = 0$). When a mother delivers her infant, the covariates for the mother-infant dyad W_{ij} are collected. These include the mother's characteristics, such as age, parity, and receiving a cesarean section (C-section), and the infant's characteristics, such as sex, weight, length, and arm circumference. Again, we assume these covariates W_{ij} are not impacted by the intervention A , even though they are measured after randomization. Finally, the infant's vital status is recorded; Y_{ij} is an indicator of infant death within 28 days. Over the course of study follow-up, we observe many such deliveries, but for the primary population of interest, we restrict to N_j pre-term births, defined as born before 37 weeks of gestation. This process is repeated for the sample size of $J = 20$ facilities.

2.2 | Counterfactuals and target causal effects

We generate counterfactual outcomes by replacing the structural equation f_A in causal model (Equation 1) with our desired intervention.³⁶ Let $Y_{ij}(a)$ be the counterfactual outcome for individual i in cluster j if, possibly contrary-to-fact, their cluster received treatment-level $A = a$. In the PTBi study, $Y_{ij}(1)$ represents the 28-day vital status for infant i in facility j if that facility had been randomized to the intervention arm ($A_j = 1$), while $Y_{ij}(0)$ represents the 28-day vital status for infant i in facility j if that facility had been randomized to the control arm ($A_j = 0$).

We can also define counterfactual outcomes at the cluster-level by taking aggregates of the individual-level ones. Many such summary measures are possible. In line with common practice,^{1,28,32-34,37,38} we focus on weighted sums of the N_j participants from cluster j :

$$Y_j^c(a) \equiv \sum_{i=1}^{N_j} \alpha_{ij} Y_{ij}(a). \quad (2)$$

Often, this weight is selected to be the inverse of the cluster size and, thus, constant across participants in a cluster: $\alpha_{ij} = 1 / N_j$ for $i = \{1, \dots, N_j\}$. With this choice, $Y_j^c(a)$ is the average counterfactual outcome for the N_j participants in cluster j . For a binary individual-level outcome, using the inverse cluster size yields a cluster-level outcome $Y^c(a)$ corresponding

to a proportion or probability. In PTBi, for example, $Y_j^c(a)$ is the counterfactual cumulative incidence of death by 28-days if, possibly contrary-to-fact, facility j received treatment-level $A_j = a$. Of course, other weighting schemes can be used to summarize the individual-level counterfactual outcomes to the cluster-level.

By applying a summary measure to the distribution of counterfactual outcomes, we define the causal effect corresponding to our research query. A wide variety of causal parameters can be expressed as the empirical mean over the sample:^{34,37-41}

$$\Phi^{c,J}(a) = \frac{1}{J} \sum_{j=1}^J \left(\sum_{i=1}^{N_j} \alpha_{ij} Y_{ij}(a) \right) = \frac{1}{J} \sum_{j=1}^J Y_j^c(a), \quad (3)$$

where α_{ij} again is the user-specified weight. Throughout, superscript c denotes a summary of cluster-level outcomes, and superscript J denotes a sample parameter. As the number of clusters grows ($J \rightarrow \infty$), the sample parameter converges to the expectation over the target population of clusters:

$$\Phi^c(a) = \mathbb{E}[Y^c(a)]. \quad (4)$$

In words, $\Phi^c(a)$ is the expected cluster-level counterfactual outcome if all clusters in the population had been assigned to treatment-level $A = a$, whereas $\Phi^{c,J}(a)$ in Equation (3) is the average cluster-level counterfactual outcome for the J clusters in the CRT. For the PTBi study and weights $\alpha_{ij} = 1 / N_j$, $\Phi^c(a)$ represents the expected incidence of 28-day mortality among preterm infants if all health facilities in the target population had been assigned to treatment arm $A = a$, while $\Phi^{c,J}(a)$ represents the average incidence if the $J = 20$ health facilities included in the study had been assigned to treatment arm $A = a$.

By taking contrasts of these treatment-specific causal parameters, we can define causal effects on any scale of interest. For example, we may be interested in the relative effect for the J clusters included in the study: $\Phi^{c,J}(1) \div \Phi^{c,J}(0)$. Alternatively, we may be interested in their difference at the population-level (a.k.a., the average treatment effect): $\Phi^c(1) - \Phi^c(0)$. For simplicity we refer to contrasts defined using Equation (3) or (4) as “cluster-level effects.”

Letting $N_T \equiv \sum_j N_j$ be the total number of participants in the CRT, we can also use Equation (3) to define causal effects at the individual-level by setting $\alpha_{ij} = J / N_T$:

$$\Phi^J(a) = \frac{1}{J} \sum_{j=1}^J \sum_{i=1}^{N_j} \left(\frac{J}{N_T} Y_{ij}(a) \right) = \frac{1}{N_T} \sum_{j=1}^J \sum_{i=1}^{N_j} Y_{ij}(a). \quad (5)$$

As sample size grows ($J \rightarrow \infty$), this converges to the expectation over the target population of clusters, each containing a finite number of participants:

$$\Phi(a) = \mathbb{E}[Y(a)]. \quad (6)$$

In words, $\Phi(a)$ is the expected individual-level outcome if all clusters in the target population received treatment-level $A = a$, whereas $\Phi^J(a)$ in Equation (5) is the average individual-level outcome for the N_j participants in the CRT. In the PTBi study, $\Phi(a)$ represents the counterfactual risk of mortality for a preterm infant if all health facilities in the target population received treatment-level $A = a$, while $\Phi^J(a)$ represents the counterfactual proportion of preterm infants who would die if the $J = 20$ health facilities included in the study had been assigned to treatment arm $A = a$. As before, we can take the difference or ratio of these treatment-specific parameters to define causal effects. For simplicity we refer to contrasts defined using Equation (5) or (6) as “individual-level effects.”

In summary, we can define a wide variety of causal parameters by considering alternative summary measures of the individual-level or cluster-level counterfactual outcomes. Further generalizations are available in Appendix A of the Data S1. When the cluster size varies (ie, $N_j \neq n, \forall j$), causal parameters giving equal weight to clusters (eg, Equation 3) will generally differ from causal parameters giving equal weight to participants (eg, Equation 5). As a toy example, suppose we have $J = 5$ clusters of varying sizes, specifically $N_j = 10$ for $j = \{1, \dots, 4\}$ and $N_5 = 10\,000$. Further suppose the counterfactual probability of the individual-level outcome depends on cluster size, such that the total number of outcomes ($\sum_i^{N_j} Y_{ij}(a)$) is 2 for clusters $j = \{1, \dots, 4\}$ and 7500 for cluster $j = 5$. Then using the inverse cluster size as weight ($\alpha_j = 1 / N_j$), the cluster-level sample parameter, given in Equation (3), would be $\Phi^{c,J}(a) = (2/10 + 2/10 + 2/10 + 2/10 + 7500/10\,000) \div 5 = 0.31$. In contrast, the individual-level sample parameter, given in Equation (5), would be $\Phi^J(a) = (2 + 2 + 2 + 2 + 7500) \div (10 + 10 + 10 + 10 + 10\,000) = 0.748$. While this is an extreme example, it illustrates the potential divergence of the treatment-specific means when cluster size varies.

Additionally since causal effects are defined through contrasts of these treatment-specific means, effects defined at the cluster-level or individual-level can also be meaningfully different when cluster size varies. The potential divergence in these causal effects is exacerbated under informative cluster size, occurring when the intervention effect is modified by cluster size.^{8,9} This scenario is explored in detail in the second simulation study and real data example, below. Nonetheless, it is worth emphasizing that even if the cluster size is constant (ie, $N_j = n, \forall j$), the cluster-level and individual-level effects have subtly different interpretations, as discussed previously. In observational studies, failing to recognize that effects defined at the aggregate-level (ie, at the cluster-level) may differ from effects defined at the individual-level is known as the “ecological fallacy.”⁴²

For completeness in the above, we defined both sample-specific and population-level measures. For ease of comparison of analytic methods, we focus on the population-level

effects (Equations 4 and 6) for the remainder of the paper and refer the reader to References 37-41,43 for further discussion on the trade-offs between targeting population vs sample effects. In brief, the sample effect might be more appealing when clusters are selected for convenience and can be estimated more precisely than the population effect.

2.3 | Observed data and identification of causal effects

For a given cluster, the observed data are the set of measured cluster-level covariates, the matrix of individual-level covariates, the randomized treatment, and the vector of individual-level outcomes:

$$O = (E, \mathbf{W}, A, \mathbf{Y}).$$

We assume the observed data are generated by sampling J times from a data generating process compatible with the above causal model (Equation 1). This provides a link between the causal model and the statistical model, which is the set of possible distributions of the observed data.¹¹ The causal model encodes that the cluster-level treatment is randomized, but does not otherwise place restrictions on the joint distribution of observed data, denoted \mathbb{P}_0 . Importantly, there are not any parametric restrictions on how the outcome vector \mathbf{Y} is generated; instead, it may be any function of the covariates (E, \mathbf{W}) , the intervention A , and unmeasured factors $U_{\mathbf{Y}}$. Altogether, our statistical model is semi-parametric.

As with the counterfactual outcome vector $\mathbf{Y}(a)$, we can summarize the observed outcome vector \mathbf{Y} in a variety of ways. We again consider weighted sums of the individual-level outcomes within a cluster:

$$Y_j^c \equiv \sum_{i=1}^{N_j} \alpha_{ij} Y_{ij}, \tag{7}$$

where α_{ij} matches the definition of the cluster-level counterfactual outcome in Equation 2. For the remainder of the paper, we focus the cluster-level outcome Y_j^c defined as the empirical mean within each cluster, corresponding to weights $\alpha_{ij} = \frac{1}{N_j}$ for $i = \{1, \dots, N_j\}$.

However, as previously discussed, we can consider alternative weighting schemes α_{ij} depending on our research question.

To identify the expected counterfactual outcome as a function of the observed data distribution and define our target statistical estimand, we require the following two assumptions. Both are satisfied by design in a CRT. First, there must be no unmeasured confounding, such that $A \perp\!\!\!\perp \mathbf{Y}(a)$. Second, there must be a positive probability of receiving each treatment-level: $\mathbb{P}_0(A = a) > 0$. Given these conditions are met in a CRT, we can express the cluster-level causal parameter $\Phi^c(a) = \mathbb{E}[Y^c(a)]$ as the expected cluster-level outcome under the treatment-level of interest $\mathbb{E}_0[Y^c | A = a]$, where subscript 0 is used to denote the observed data distribution \mathbb{P}_0 ; a proof is provided in Reference 28. Likewise, the individual-

level causal parameter $\Phi(a) = \mathbb{E}[Y(a)]$ equals the expected individual-level outcome under treatment-level of interest: $\mathbb{E}_0[Y | A = a]$, recognizing the slight abuse to notation because the observed data $O \sim \mathbb{P}_0$ are at the cluster-level.

We can gain efficiency in CRTs by adjusting for baseline covariates (eg, References 1,7,29, and 44). Specifically, the treatment-specific expectation of the cluster-level outcome $\mathbb{E}_0[Y^c | A = a]$ can be expressed as the conditional expectation of the cluster-level outcome Y^c given the treatment-level of interest a and the baseline covariates (E, \mathbf{W}) , averaged over the covariate distribution: $\mathbb{E}_0[\mathbb{E}_0(Y^c | A = a, E, \mathbf{W})]$. In practice, we cannot directly adjust for the entire matrix of individual-level covariates \mathbf{W} during estimation. However, we can still improve precision by including lower-dimensional summary measures of \mathbf{W} in the adjustment set or by aggregating the individual-level conditional mean outcome to the cluster-level (details in Appendix B of the Data S1).²⁸ For simplicity, we use W^c to denote either approach to including individual-level covariates in our cluster-level statistical estimand, defined as

$$\Psi_0^c(a) \equiv \mathbb{E}_0[\mathbb{E}_0(Y^c | A = a, E, W^c)]. \quad (8)$$

As previously discussed, under the randomization and positivity assumptions, both holding by design, this equals the treatment-specific mean of the cluster-level counterfactual outcomes $\Phi^c(a) = \mathbb{E}[Y^c(a)]$.

Likewise, our statistical estimand for the treatment-specific mean of the individual-level counterfactual outcomes $\Phi(a) = \mathbb{E}[Y(a)]$ is

$$\Psi_0(a) \equiv \mathbb{E}_0[\mathbb{E}_0(Y | A = a, E, W)], \quad (9)$$

again acknowledging the slight abuse to notation, because subscript 0 denotes the distribution of the cluster-level data $O \sim \mathbb{P}_0$. Our individual-level estimand $\Psi_0(a)$ is the conditional expectation of the individual-level outcome Y given the treatment-level of interest a , the cluster-level covariates E and the individual-specific covariates W , averaged over the covariate distribution. For both estimands (Equations 8 and 9), we emphasize that covariate adjustment is being used for efficiency gains only and not to control for confounding or other sources of bias, such as missing outcomes or selection.^{16-18,20-23}

As before, we take contrasts of the cluster-level or individual-level estimands corresponding to our causal effect of interest. To give context for the methods comparison, we will focus on the relative scale for the remainder of the paper; however, our discussion is equally applicable to other scales (eg, additive or odds ratio). Specifically, the relative effect is identified at the cluster-level as

$$\frac{\Psi_0^c(1)}{\Psi_0^c(0)} = \frac{\mathbb{E}_0[\mathbb{E}_0(Y^c | A = 1, E, W^c)]}{\mathbb{E}_0[\mathbb{E}_0(Y^c | A = 0, E, W^c)]}, \quad (10)$$

and at the individual-level as

$$\frac{\Psi_0(1)}{\Psi_0(0)} = \frac{\mathbb{E}_0[\mathbb{E}_0(Y | A = 1, E, W)]}{\mathbb{E}_0[\mathbb{E}_0(Y | A = 0, E, W)]}. \quad (11)$$

As detailed below, most analytic methods naturally only estimate one of the above statistical estimands, whereas few have the flexibility to estimate both. While we focused on identification of population-level effects, the extensions to sample effects are fairly straightforward, as discussed in Reference 41.

3 | STATISTICAL ESTIMATION AND INFERENCE

In this section, we compare the methods commonly used to analyze CRTs. We describe their target of inference and their ability to adjust for baseline covariates to improve precision and thereby improve statistical power. We broadly consider two classes of estimation methods: approaches using only cluster-level data and approaches using both individual- and cluster-level data. The former immediately aggregate the data to the cluster-level and can only adjust for cluster-level covariates, while the latter allow for adjustment of individual-level covariates, an appealing option, as these pair naturally with individual-level outcomes. Examples of cluster-level approaches include the t-test and cluster-level TMLE. Cluster-level approaches naturally target cluster-level effects (eg, Equation 10), but with the appropriate choice of weights can also target individual-level effects (eg, Equation 11). Examples of approaches using individual-level data include GEE, CARE, and Hierarchical TMLE^{1,26,28}. These individual-level approaches often estimate different causal effects, as detailed below. To the best of our understanding, of the algorithms discussed here, Hierarchical TMLE is the only individual-level approach that can estimate effects defined at the cluster-level (eg, Equation 10) or at the individual-level (eg, Equation 11). It should be possible to modify G-computation and inverse probability weighting to target both cluster-level and individual-level effects, but these extensions are beyond the scope of this paper. It is not well understood how to incorporate weights for GEE.⁴⁵

We now define the notation used throughout this section. Recall the cluster level-outcome Y^c is defined as a weighted sum of individual-level outcomes, as in Equation (7). We denote the conditional expectation of the cluster-level outcome Y^c given the cluster-level intervention $A = a$ and covariates (E, W^c) as

$$\mu^c(a, E, W^c) \equiv \mathbb{E}(Y^c | A = a, E, W^c). \quad (12)$$

Likewise, we denote the conditional expectation of the individual-level outcome Y given the cluster-level intervention $A = a$, the cluster-level covariates E , and that individual's covariates W as

$$\mu(a, E, W) \equiv \mathbb{E}(Y \mid A = a, E, W). \quad (13)$$

Throughout, we refer to Equation (12) and to (13) as the cluster-level and individual-level outcome regressions, respectively. The unadjusted expectations of the cluster-level and individual-level outcomes within treatment arm a are defined as $\mu^c(a) \equiv \mathbb{E}(Y^c \mid A = a)$ and $\mu(a) \equiv \mathbb{E}(Y \mid A = a)$, respectively. We denote the cluster-level propensity score as

$$\pi^c(a \mid E, W^c) \equiv \mathbb{P}(A = a \mid E, W^c), \quad (14)$$

and the individual-level propensity score as

$$\pi(a \mid E, W) \equiv \mathbb{P}(A = a \mid E, W). \quad (15)$$

We define unadjusted probabilities $\pi^c(a)$ and $\pi(a)$, analogously.

3.1 | Analytic approaches using cluster-level data

Cluster-level approaches obtain point estimates and inference after the individual-level data have been aggregated to the cluster-level.^{1,7} Most commonly, this aggregation is done by taking the empirical mean within each cluster. However, as previously detailed in Section 2.2, we can consider several ways to summarize the individual-level data to the cluster-level (ie, different α_{ij} weighting schemes).

3.1.1 | Unadjusted effect estimator—Once the data are aggregated to the cluster-level, a common approach for estimation and inference is based on contrasts of the treatment-specific average outcomes:

$$\hat{\mu}^c(a) = \frac{1}{J} \sum_{j=1}^J \frac{\mathbb{1}(A_j = a) Y_j^c}{\hat{\pi}^c(a)}, \quad (16)$$

where $\hat{\pi}^c(a)$ denotes the unadjusted estimate of the cluster-level propensity score (ie, the proportion of clusters in the trial receiving treatment-level $A = a$). For simplicity for the remainder of the manuscript, we assume the trial has equal allocation of arms, such that $\hat{\pi}^c(a) = 1/2$; however, our results should generalize to trials with more than two arms and to imbalanced trials. Then, if we let $Y_{a,k}^c$ denote the cluster-level outcome for observation $k = \{1, \dots, J/2\}$ in treatment arm $A = a$, the treatment-specific mean simplifies

to $\hat{\mu}^c(a) = \frac{1}{J/2} \sum_{k=1}^{J/2} Y_{a,k}^c$. In PTBi, $\hat{\mu}^c(a)$ represents the average incidence of 28-day infant mortality among facilities that received treatment-level $A = a$. We obtain a point estimate of the cluster-level effect by contrasting $\hat{\mu}^c(1)$ and $\hat{\mu}^c(0)$ on the scale of interest and obtain statistical inference using the t -distribution. Suppose, for example, we were interested in the cluster-level average treatment effect $\mathbb{E}[Y^c(1)] - \mathbb{E}[Y^c(0)]$; then our point estimate would be $\hat{\mu}^c(1) - \hat{\mu}^c(0)$ and we would test the null hypothesis using a Student's t -test. Statistical power may be improved by considering alternative weighting schemes when summarizing individual-level outcomes to the cluster-level;¹ however, as previously discussed, different weights α_{ij} imply different target effects.

For the relative effect, applying the logarithmic transformation is sometimes recommended when the cluster-level summaries are skewed, which may be more common for rate-type outcomes.¹ However, it is important to note that depending on how this transformation is implemented, estimation and inference may be for the ratio of the geometric means, as opposed to the ratio of the arithmetic means. (Recall for J observations of some variable X , the geometric mean is $(\prod_{i=1}^J X_i)^{1/J}$, whereas the arithmetic mean is $1/J \sum_{i=1}^J X_i$.) Specifically, suppose we first take the log of the cluster-level outcomes and then take the average within each arm:

$$\bar{I}_a \equiv \frac{1}{J/2} \sum_{k=1}^{J/2} \log(Y_{a,k}^c) = \log \left(\prod_{k=1}^{J/2} Y_{a,k}^c \right)^{\frac{1}{J/2}}, \quad (17)$$

where again $Y_{a,k}^c$ denotes the cluster-level outcome for cluster $k = \{1, \dots, J/2\}$ in treatment arm a . Applying a t -test to the difference in these treatment-specific means $\bar{I}_1 - \bar{I}_0$ (and then exponentiating) targets the ratio of the geometric means. Continuing our toy example from Section 2.2, the arithmetic mean of the cluster-level outcomes in the control was $\Phi^{c,J}(0) = 0.31$, while the geometric mean of the cluster-level control outcomes would be 0.26. To avoid changing the target of inference, we can instead apply the Delta Method to obtain point estimates and inference for the ratio of arithmetic means, as in Equation (10). We refer the reader to Reference 46 for more details.

3.1.2 | Cluster-level TMLE with adaptive prespecification—As previously discussed, statistical power is often improved by adjusting for baseline covariates that are predictive of the outcome. Once the data have been aggregated to the cluster-level, we can proceed with estimation and inference for the cluster-level estimand $\Psi_0^c(a) = \mathbb{E}_0[\mathbb{E}_0(Y^c | A = a, E, W^c)]$, using methods for i.i.d. data. Examples of common algorithms for $\Psi_0^c(a)$ include parametric G-computation, inverse probability of treatment weighting estimators (IPTW), and TMLE.¹¹ Due to treatment randomization, these algorithms will be consistent, even under misspecification of the outcome regression.⁴⁷ Given that G-computation and IPTW are well-established approaches, we focus on TMLE, which is a general class of double robust, semiparametric efficient, plug-in estimators.¹¹

Here, we briefly review the steps of a cluster-level TMLE and then present a solution for optimal selection of the adjustment covariates in trials with limited numbers of clusters. We conclude with a discussion of how to apply weights to the cluster-level TMLE to estimate an individual-level estimand (eg, $\Psi_0(a) = \mathbb{E}_0[\mathbb{E}_0(Y | A = a, E, w)]$). In the next section, we discuss an alternative implementation of TMLE that can harness both individual-level and cluster-level covariates to increase precision, while maintaining Type-I error control.

To implement a cluster-level TMLE of the cluster-level effect, we first obtain an initial estimator of the expected cluster-level outcome $\mu^c(A, E, W^c)$. Next, we update this initial estimator $\hat{\mu}^c(A, E, W^c)$ using information contained in the estimated propensity score $\hat{\pi}^c(a | E, W^c)$. Specifically, we define the ‘‘clever covariate’’ as the inverse of the estimated propensity score for cluster j :

$$\hat{H}^c(a, E_j, W_j^c) = \frac{\mathbb{1}(A_j = a)}{\hat{\pi}^c(a | E_j, W_j^c)}.$$

Then on the logit-scale, we regress the cluster-level outcome Y_j^c on the covariates $\hat{H}^c(1, E_j, W_j^c)$ and $\hat{H}^c(0, E_j, W_j^c)$ with the initial estimator $\hat{\mu}^c(A_j, E_j, W_j^c)$ as the offset. This provides the following targeted estimator, while simultaneously solving the efficient score equation:

$$\hat{\mu}^{c*}(a, E, W^c) = \text{logit}^{-1}[\text{logit}(\hat{\mu}^c(a, E, W^c)) + \hat{\epsilon}_1 \hat{H}^c(1, E, W^c) + \hat{\epsilon}_0 \hat{H}^c(0, E, W^c)], \quad (18)$$

where $\hat{\epsilon}_1$ and $\hat{\epsilon}_0$ denote the estimated coefficients for $\hat{H}^c(1, E, W^c)$ and $\hat{H}^c(0, E, W^c)$, respectively. Finally, we obtain a point estimate of the treatment-specific mean $\Psi_0^c(a)$ by averaging the targeted predictions of the cluster-level outcomes across the J clusters:

$$\hat{\Psi}^{c*}(a) = \frac{1}{J} \sum_{j=1}^J \hat{\mu}^{c*}(a, E_j, W_j^c).$$

To evaluate the intervention effect, we contrast our estimates $\hat{\Psi}^{c*}(1)$ and $\hat{\Psi}^{c*}(0)$ on the scale of interest and apply the Delta Method for inference. The variance of asymptotically linear estimators, such as the TMLE, may be estimated using the estimator’s influence function.¹¹ These types of estimators enjoy properties that follow from the Central Limit Theorem, allowing us to construct 95% Wald-type confidence intervals. As a finite sample approximation to the normal distribution, we recommend using the t -distribution with $J - 2$ degrees of freedom. For the treatment-specific mean $\Psi_0^c(a)$, for example, the influence function and asymptotic variance for the cluster-level TMLE are well-approximated as $\hat{D}^c(a) = \hat{H}^c(a, E, W^c) \times (Y^c - \hat{\mu}^{c*}(a, E, W^c)) + \hat{\mu}^{c*}(a, E, W^c) - \hat{\Psi}^{c*}(a)$ and $\widehat{\text{Var}}[\hat{D}^c(a)] / J$, respectively.

In a CRT, the propensity score $\pi^c(a, E, W^c) = \pi^c(a)$ is known and does not need to be estimated. However, further gains in efficiency can be achieved through estimation of the propensity score.^{44,47} If both the cluster-level outcome regression and cluster-level propensity score are consistently estimated, the TMLE will be an asymptotically efficient estimator. However, consistent estimation of the outcome regression is nearly impossible when using an a priori-specified regression model.

To improve precision while preserving Type-I error control, we previously proposed “Adaptive Prespecification,” a supervised learning approach using sample-splitting to choose the adjustment set that maximizes efficiency.²⁹ In brief, we prespecify a set of candidate generalized linear models for the cluster-level outcome regression $\mu^c(a, E, W^c)$ and propensity score $\pi^c(a | E, W^c)$. To avoid forced adjustment at the detriment of precision, the unadjusted estimator should always be included as a candidate. We also prespecify a cross-validation scheme; for small trials (eg, $J \leq 30$), we recommend leave-one-cluster-out. To measure performance, we prespecify the squared influence function as our loss function. Then we choose the candidate estimator of $\mu^c(a, E, W^c)$ that minimizes the cross-validated variance estimate using the influence function based on the known propensity score (ie, $\pi^c(a) = 0.5$). We then select the candidate estimator of propensity score $\pi^c(a | E, W^c)$ that further minimizes the cross-validated variance estimate using the influence function when combined with the previously selected estimator $\hat{\mu}^c(a, E, W^c)$. Together, the selected estimators $\hat{\mu}^c(a, E, W^c)$ and $\hat{\pi}^c(a | E, W^c)$ form the “optimal” TMLE according to the principle of empirical efficiency maximization.⁴⁴

Application of TMLE with Adaptive Prespecification to real data from the SEARCH CRT resulted in notable precision gains.^{16,48} As compared to the unadjusted estimator, TMLE was 4.6 times more efficient for the effect on HIV incidence, 2.6 times more efficient for the effect on the incidence tuberculosis, and 1.8 times more efficient for the effect on hypertension control. Given that this cluster-level TMLE adjusted for at most two cluster-level variables, these gains in efficiency may seem surprising. However, we believe that they demonstrate the power of using Adaptive Prespecification to flexibly select the optimal adjustment strategy to maximize empirical efficiency. These gains are also seen in a recent extension of Adaptive Prespecification for flexible adjustment of many covariates in trials with a larger number of randomized units.⁴⁹

3.1.3 | Targeting an individual-level effect with an estimator using cluster-level data—With or without Adaptive Prespecification, the cluster-level TMLE naturally estimates cluster-level parameters (eg, Equation (10)) corresponding to contrasts of cluster-level counterfactuals, such as $\mathbb{E}[Y^c(1)] \div \mathbb{E}[Y^c(0)]$. However, the cluster-level TMLE can also estimate individual-level parameters (eg, Equation 11) corresponding to contrasts of individual-level counterfactuals, such as $\mathbb{E}[Y(1)] \div \mathbb{E}[Y(0)]$. To do so, we include weights $J / N_T \times 1 / \alpha_j$ in each step of the cluster-level analysis (derivation and R code in the Data S1). This approach may be relevant when data are only available at the cluster-level, but interest is in an individual-level effect. Finally, we note that the unadjusted effect estimator can be considered a special case of the TMLE where the adjustment set is empty:

$(E, W^c) = \{\emptyset\}$. Therefore, the unadjusted effect estimator based on cluster-level data can also estimate an individual-level effect if the appropriate weights are applied.

3.2 | Analytic approaches using individual-level data

We now discuss how to leverage individual-level covariates when estimating effects in CRTs. This is done by aggregating to the cluster-level *after* estimating the expected individual-level outcome or by implementing a fully individual-level approach. In all cases, clustering must be accounted for during variance estimation. Estimators using individual-level data vary in their flexibility to estimate both cluster-level effects (eg, Equation (10)) and individual-level effects (eg, Equation 11).

3.2.1 | Hierarchical TMLE—In Section 3.1, we discussed a cluster-level TMLE for estimating effects in CRTs based on aggregating the data to the cluster-level. An alternative and equally valid approach for CRT analysis is to ignore clustering when obtaining a point estimate and then account for clustering during variance estimation and statistical inference.^{50,51} Such an approach naturally estimates an individual-level parameter, such as $\Psi_0(a) = \mathbb{E}_0[\mathbb{E}_0(Y | A = a, E, W)]$. We now present an individual-level TMLE using this alternative approach and refer to it as “Hierarchical TMLE” to emphasize its distinction from the standard TMLE for an individually randomized trial with N_T i.i.d. participants. In Appendix B, we also present a “Hybrid TMLE” that obtains an initial estimate of the cluster-level outcome regression based on aggregates of the individual-level outcome regression (ie, $\hat{\mu}^c(A_j, E_j, W_j^c) = \sum_i^{N_j} \alpha_{ij} \hat{\mu}(A_j, E_j, W_{ij})$) and then proceeds with estimation and inference as outlined in Section 3.1.2. Both approaches leverage the natural pairing of individual-level outcomes with individual-level covariates, extend to a pair-matched design,²⁸ and can be combined with weights to estimate effects at the cluster-level or individual-level.

To implement Hierarchical TMLE for the individual-level effect, we pool participant-level data across clusters to obtain estimators of the individual-level outcome regression $\mu(A, E, W)$ and the individual-level propensity score $\pi(a | E, W)$. The initial outcome regression estimator $\hat{\mu}(A, E, W)$ is then updated based on the estimated propensity score $\hat{\pi}(a | E, W)$. As before, we calculate the “clever covariate,” but now at the individual-level:

$$\hat{H}(a, E_j, W_{ij}) = \frac{\mathbb{1}(A_j = a)}{\hat{\pi}(a | E_j, W_{ij})},$$

for $a = \{1, 0\}$. Then on the logit-scale, we regress the individual-level outcome Y_{ij} on the individual-level covariates $\hat{H}(1, E_j, W_{ij})$ and $\hat{H}(0, E_j, W_{ij})$ with the initial individual-level estimator $\hat{\mu}(A_j, E_j, W_{ij})$ as the offset. This provides the following updated estimator of the expectation of the individual-level outcome, while simultaneously solving the efficient score equation:

$$\hat{\mu}^*(a, E, W) = \text{logit}^{-1}[\text{logit}(\hat{\mu}(a, E, W)) + \hat{\epsilon}_1 \hat{H}(1, E, W) + \hat{\epsilon}_0 \hat{H}(0, E, W)], \quad (19)$$

where $\hat{\epsilon}_1$ and $\hat{\epsilon}_0$ now denote the estimated coefficients for $\widehat{H}(1, E, W)$ and $\widehat{H}(0, E, W)$. Then we obtain a point estimate of the treatment-specific mean $\Psi_0(a)$ by averaging these targeted predictions:

$$\widehat{\Psi}^*(a) = \frac{1}{N_T} \sum_{j=1}^J \sum_{i=1}^{N_j} \hat{\mu}^*(a, E_j, W_{ij}),$$

where, again, N_T denotes the total number of participants across all clusters. Thus, this implementation of Hierarchical TMLE to estimate an individual-level parameter is nearly identical to standard TMLE for i.i.d. data. Key distinctions are in sample-splitting (if used) and variance estimation, both of which must respect the cluster as the independent unit in CRTs. Specifically, to estimate the variance of Hierarchical TMLE for $\widehat{\Psi}^*(a)$, we aggregate an individual-level influence function to the cluster-level and then take the sample variance of the estimated cluster-level influence function, scaled by the number of independent units J . For example, the influence function for this TMLE of $\widehat{\Psi}^*(a)$ is well-approximated by $\widehat{D}_j(a) = \sum_{i=1}^{N_j} J / N_T \times \widehat{D}_{ij}(a)$ where $\widehat{D}_{ij}(a) = \widehat{H}(a, E_j, W_{ij}) \times (Y_{ij} - \hat{\mu}^*(a, E_j, W_{ij})) + \hat{\mu}^*(a, E_j, W_{ij}) - \widehat{\Psi}^*(a)$. (See Schnitzer et al for a proof.⁵¹) Altogether, this implementation of Hierarchical TMLE naturally estimates individual-level effects (eg, Equation 11) and is analogous to using an independent working correlation matrix with the robust variance estimator in GEE,⁴⁵ described below.

Unlike GEE, however, Hierarchical TMLE can also easily estimate cluster-level effects (eg, Equation 10).²⁸ To do so, we incorporate the α_{ij} weights throughout the analysis to obtain cluster-level point estimates and inference. Specifically, when obtaining a point estimate, we aggregate the targeted predictions within clusters before averaging across clusters: $\widehat{\Psi}^{c*}(a) = 1 / J \sum_j \sum_i \alpha_{ij} \hat{\mu}^*(a, E_j, W_{ij})$. Likewise, we estimate a cluster-level influence function for this TMLE of $\widehat{\Psi}^{c*}(a)$ as $\widehat{D}_j(a) = \sum_{i=1}^{N_j} \alpha_{ij} \times \widehat{D}_{ij}(a)$ where $\widehat{D}_{ij}(a) = \widehat{H}(a, E_j, W_{ij}) \times (Y_{ij} - \hat{\mu}^*(a, E_j, W_{ij})) + \hat{\mu}^*(a, E_j, W_{ij}) - \widehat{\Psi}^{c*}(a)$. Further details are available in Balzer et al²⁸ and R code is provided in the Data S1.

For either the cluster-level or individual-level effect, the variance of Hierarchical TMLE is well-approximated by the variance of the cluster-level influence function, scaled by the number of independent units J . As before, the Delta Method is applied to obtain point estimates and inference for the intervention effect on any scale of interest. Again, we recommend the t-distribution with $J - 2$ degrees of freedom for confidence interval construction and testing the null hypothesis. Hierarchical TMLE will be an asymptotically efficient estimator if the outcome regression and propensity score are consistently estimated at reasonable rates. In practice, we again recommend using Adaptive Prespecification to select the optimal adjustment strategy to maximize efficiency.

3.2.2 | Covariate-adjusted residuals estimator—The covariate-adjusted residuals estimator (CARE) was first proposed in Gail et al²⁷ and later popularized by Hayes and Moulton.¹ CARE is implemented by pooling individual-level data across clusters and then

regressing the individual-level outcome Y on the individual-level and cluster-level covariates of interest (W, E), but not the cluster-level intervention A . Then the predictions from this regression are aggregated to the cluster-level. Finally, a t -test comparing the mean residuals (ie, the discrepancies between observed and predicted outcomes) by arms is performed, since the average residuals should be the same between arms under the null hypothesis. CARE naturally targets cluster-level effects, as in Equation (10), and it is not immediately obvious how to incorporate weights to target individual-level effects, as in Equation (11).

As a concrete example, suppose our goal is to estimate the cluster-level effect on the relative scale and the individual-level outcome Y is binary. To implement CARE in this setting, we first fit an individual-level logistic regression, such as

$$\mathbb{E}(Y | E, W) = \text{logit}^{-1}[\alpha + \beta_E E + \beta_W W], \quad (20)$$

where β_E and β_W denote the magnitude by which the log odds of the outcome for the i th individual in the j th cluster is affected (linearly) by the cluster-level covariates E_j and individual-level covariates W_{ij} , respectively. From this regression we obtain the expected number of events in the j th cluster as $e_j = \sum_i^{N_j} \text{logit}^{-1}(\hat{\alpha} + \hat{\beta}_E E_j + \hat{\beta}_W W_{ij})$ and compare it with the observed number of events $d_j = \sum_i^{N_j} Y_{ij}$ through ratio-residuals:

$$R_j = \frac{d_j}{e_j}.$$

Hayes and Moulton¹ note that these ratio-residuals are often right-skewed and recommend a logarithmic transformation. Specifically, they recommend applying a t -test to obtain point estimates and inference for the difference in the treatment-specific averages of the log-transformed residuals:

$$\frac{1}{J/2} \sum_{k=1}^{J/2} \log(R_{1,k}) - \frac{1}{J/2} \sum_{k=1}^{J/2} \log(R_{0,k}) = \log \left(\prod_{k=1}^{J/2} R_{1,k} \right)^{\frac{1}{J/2}} - \log \left(\prod_{k=1}^{J/2} R_{0,k} \right)^{\frac{1}{J/2}}, \quad (21)$$

where $R_{a,k}$ denotes the ratio-residual for cluster $k = \{1, \dots, J/2\}$ in arm a . As detailed in Section 3.1.1, after exponentiation, we recover estimates and inference for the ratio of the geometric means and thereby a different causal effect than the standard risk ratio, given in Equation (10). A straightforward extension to pair-matched design is illustrated in Reference 1.

3.2.3 | Generalized estimating equations—We now consider a class of estimating equations, sometimes referred to as “population-average models,” for estimating effects in CRTs.⁵⁰ In GEE, estimation and inference is conducted at the individual-level and a working correlation matrix is used to account for the dependence of outcomes within clusters. GEE

naturally targets individual-level effects, as in Equation (11). As with CARE, it is unclear whether weights can be incorporated to instead target cluster-level effects, as in Equation (10). Indeed, recent work by Wang et al⁴⁵ suggest that when targeting the cluster-level effect, inappropriate use of weights in GEE can result in meaningful bias.

In GEE, the expected individual-level outcome is modeled a function of the treatment and possibly covariates of interest.^{50,52} Specifically, consider the following “marginal model” for the expected individual-level outcome $\mathbb{E}(Y | A)$:

$$\mu(A_j) = g^{-1}(\beta_0 + \beta_A A_j), \quad (22)$$

where $g^{-1}(\cdot)$ denotes the inverse-link function. Commonly, the identity link is used for continuous outcomes, the log-link for count outcomes, and the logit-link for binary outcomes. Effect estimation in GEE is usually done by obtaining a point estimate and inference for the treatment coefficient β_A . Then at the individual-level, β_A represents the additive causal effect for the identity link; e^{β_A} represents the relative effect for the log-link, and e^{β_A} represents the odds ratio effect for the logit-link. In other words, the link function often determines the scale on which the effect is estimated.

As with other CRT approaches, GEE may improve efficiency by adjusting for covariates. Consider, for example, the following “conditional model” for the expected individual-level outcome $\mathbb{E}(Y | A, E, W)$:

$$\mu(A_j, E_j, W_{ij}) = g^{-1}(\beta_0 + \beta_A A_j + \beta_E E_j + \beta_W W_{ij}), \quad (23)$$

where again $g^{-1}(\cdot)$ denotes the inverse-link function. Except for linear and log-linear models without interaction terms, the interpretation of β_A is generally not the same as in the marginal model⁵² For the logistic link function, for example, β_A in Equation (23) would yield the conditional log-odds ratio, instead of the marginal log-odds ratio. However, a recent modification to GEE, presented in the next subsection, allows for estimation of marginal effects, while adjusting for individual-level or cluster-level covariates.

For either a marginal or conditional specification, the GEE estimator solves the following equation:

$$\sum_{j=1}^J \mathbf{D}_j^T \mathbf{V}_j^{-1} (\mathbf{Y}_j - \boldsymbol{\mu}_j) = 0,$$

where $\boldsymbol{\mu}_j$ is the vector containing individual-level outcome regressions for cluster j , $\mathbf{D}_j = \frac{\partial \boldsymbol{\mu}_j}{\partial \boldsymbol{\beta}}$ is the gradient matrix, and \mathbf{V}_j is the working correlation matrix used to account for dependence of individuals within a cluster j .⁵⁰ GEE yields a consistent point estimate of

β_A under the marginal model, even if the correlation matrix has been misspecified. It is worth noting, however, that the interpretation of the estimated effect changes subtly when using alternative correlation matrices.⁵² Additionally, under misspecification of the correlation matrix, the usual SEs obtained are not valid, and the sandwich variance estimator must be used.⁵⁰ In general, unless the number of clusters is relatively large and the number of participants within cluster is relatively small, the sandwich-based SEs can underestimate the true variance of β and yield confidence intervals with coverage probability below desired nominal level. Fortunately, several corrections to variance estimation exist.^{7,53,54} Finally, for a CRT where the intervention is randomized within matched pairs of clusters, a pair-matched analysis can be conducted by specifying a fixed effect for the pair, while maintaining the correlation structure at the cluster-level. However, pair-matched analyses are generally discouraged for studies with fewer than 40 clusters.^{1,7,53}

3.2.4 | Augmented-GEE—As discussed in the previous subsection, the treatment coefficient β_A resulting from GEE does not always correspond to the marginal effect. Recently, a modification to GEE was proposed to ensure β_A can be interpreted as a marginal effect, while simultaneously adjusting for baseline covariates to improve efficiency.³⁰ This approach is referred to as Augmented-GEE (Aug-GEE) and naturally targets individual-level effects, as in Equation (11). As with standard GEE, it is unclear if weights can be incorporated in Aug-GEE to instead target cluster-level effects, as in Equation (10). Additionally, as with standard GEE, the link function $g(\cdot)$ determines the scale on which the effect is estimated in Aug-GEE (ie, additive for the identity link, relative for the log-link, and odds ratio for the logit-link).

As commonly implemented, Aug-GEE modifies GEE by including an additional “augmentation” term, which incorporates the conditional expectation of the individual-level outcome $\mu(A_j, E_j, W_{ij})$.

The general form of the Aug-GEE for a binary treatment is given by

$$\sum_{j=1}^J \left[\mathbf{D}_j^T \mathbf{V}_j^{-1} (\mathbf{Y}_j - \boldsymbol{\mu}_j(A_j)) - \sum_{a=0}^1 [\mathbb{1}(A_j = a) - \pi^c(a)] \gamma_a \right] = 0, \quad (24)$$

with augmentation term

$$\gamma_a = \mathbf{D}_j^T \mathbf{V}_j^{-1} (\boldsymbol{\mu}_j(a, E_j, W_{ij}) - \boldsymbol{\mu}_j(a)).$$

For cluster j , $\boldsymbol{\mu}_j(A_j)$ denotes the vector of marginal regressions (as in Equation 22) and $\boldsymbol{\mu}_j(A_j, E_j, W_{ij})$ denotes the vector of conditional regressions (as in Equation 23). The cluster-level propensity score, defined in Equation (14), is treated as known (ie, $\pi^c(a) = 0.5$). By solving Equation (24), we can obtain point estimates and inference for the marginal effects. As with the other estimators considered, if the conditional outcome regression $\mu(A_j, E_j, W_{ij})$ is misspecified, the resulting estimator is asymptotically normal and consistent; however,

it is not efficient.³⁰ Indeed, the efficiency of Aug-GEE heavily depends on the matrix $\mathbf{D}^T \mathbf{V}^{-1}$. Stephens et al⁵⁵ show how to further improve Aug-GEE by deriving the semi-parametric locally efficient estimator, but ultimately conclude that the high-dimensional inverse covariance matrix, required in their approach, presents a substantial barrier to any practical gains.

4 | SIMULATION STUDIES

To examine the finite sample properties of the previously discussed CRT estimators and to demonstrate the practical impact of targeting different causal effects, we conducted two simulation studies. Full R code is provided in the Data S1.

4.1 | Simulation I

We simulated a simplified data generating process reflecting the hierarchical data structure of the PTBi study, which randomized $J = 20$ clusters, corresponding to health facilities. For each cluster $j = \{1, \dots, 20\}$, we generated cluster-level covariates $E1_j \sim \text{Norm}(2, 1)$ and $E2_j \sim \text{Norm}(0, 1)$ and cluster size $N_j \sim \text{Norm}(150, 80)$ subject to a minimum of 30 participants. For each cluster j , we also simulated random variables $U_{E2_j} \sim \text{Unif}(-0.2, 1.5)$ and $U_{E2_j} \sim \text{Unif}(-0.5, 0.5)$ to act as an unmeasured source of dependence within each cluster. Then for participant $i = \{1, \dots, N_j\}$ in each cluster j , we generated four individual-level covariates: $W1_{ij} \sim \text{Norm}(2U_{E1_j}, 0.35)$; $W2_{ij} \sim \text{Norm}(4U_{E1_j}, 0.9)$; $W3_{ij} \sim \text{Norm}(U_{E2_j}, 0.5)$, and $W4_{ij} \sim \text{Norm}(U_{E2_j}, 0.5)$.

To reflect the PTBi study design, we paired clusters on $E2$ using the nonbipartite matching algorithm.⁵⁶ Within the pair, one cluster was randomized to the intervention arm ($A_j = 1$) and the other to the control arm ($A_j = 0$). Lastly, we generated the individual-level outcomes as a function of the intervention A_j , the cluster-level and individual-level covariates, and unmeasured factor $U_{Y_{ij}} \sim \text{Unif}(0, 1)$:

$$Y_{ij} \sim 1 \left[U_{Y_{ij}} < \text{logit}^{-1}(-0.75 - 0.35A_j + 0.8W1_{ij} + 0.4W2_{ij} - 0.3E1_j - 0.2A_jW2_{ij}) \right]. \quad (25)$$

To assess Type-I error control, we generated the outcomes after setting the terms involving the treatment A to 0. For a population of 2500 clusters, we also generated counterfactual outcomes under the intervention $Y(1)$ and under the control $Y(0)$ by setting $A = 1$ and $A = 0$, respectively. These counterfactuals were used to calculate the true values of the causal parameters, defined at both the cluster-level and individual-level. Specifically, we calculated the cluster-level relative effect (Equation 10), the individual-level relative effect (Equation 11), as well as the geometric incidence ratio (ie, the ratio of the geometric means of the cluster-level outcomes).

For estimation of the cluster-level relative effect, we implemented Hierarchical TMLE using Adaptive Prespecification to select from $\{\emptyset, W1, W2, W3, W4\}$ and the cluster-level TMLE with Adaptive Prespecification to select the optimal adjustment set from

$\{\emptyset, W1^c, W2^c, W3^c, W4^c\}$, where $W1^c, \dots, W4^c$ were the empirical means of their individual-level counterparts. For comparison, we also implemented a cluster-level TMLE not adjusting for any covariates, hereafter called the “unadjusted estimator.” Inference for the unadjusted estimator and the TMLEs was based on their influence functions.

CARE, GEE, and Aug-GEE relied on a fixed specification of the individual-level outcome regression with main terms adjustment for both cluster-level and individual-level covariates: $\{W1, W2, W3, W4, E1, E2\}$. For CARE, we used the logit-link to obtain outcome predictions in the absence of the treatment, then applied the log-transformation to the ratio residuals, as recommended by Hayes and Moulton,¹ and finally obtained inference with a t -test. For comparison, we also implemented a standard Student’s t -test after log-transforming the cluster-level outcomes. In GEE and Aug-GEE, we used the log-link function and Fay and Graubard’s finite sample variance correction Reference 54, as implemented in the `geesmv` and `CRTgeeDR` packages, respectively.^{57,58} All algorithms ignored the matched pairs used for intervention randomization.

4.2 | Results of simulation I

In these simulations, the cluster-level incidence ratio (Equation 10), targeted by the unadjusted estimator and the TMLEs, was identical to the individual-level risk ratio (Equation 11), targeted by GEE and Aug-GEE: 0.83. The ratio of geometric means, targeted by the t -test and CARE, was 0.81, indicating a slightly larger effect. In other words, the simulated intervention resulted in a relative reduction of the mean cluster-level outcomes of 17% on the arithmetic scale and 19% on the geometric scale. As demonstrated in the second simulation study, there can be substantial divergence between the causal parameters, and it is essential to prespecify a target effect corresponding to the research query.

While the CRT estimators targeted different effects, it is still valid to compare their attained power, defined as the proportion of times the false null hypothesis was rejected at the 5% significance level, and confidence interval coverage, defined the proportion of times the 95% confidence intervals contained the true value of the target effect. Additionally, we examined Type-I error, defined as the proportion of times the true null hypothesis was rejected at the 5% significance level. These metrics are shown in Table 1 for each estimator (for its corresponding target estimand) across 500 iterations of the data generating process, each with $J = 20$ clusters.

As shown in Table 1, the unadjusted estimator of the cluster-level relative effect achieved low statistical power (18%) with slightly conservative confidence interval coverage (97%) when there was an effect and controlled Type-I error (4%) under the null. In this simulation, the cluster-level TMLE and Hierarchical TMLE performed similarly and provided substantial efficiency gains over the unadjusted estimator of the same effect. Specifically, the TMLEs achieved a statistical power of 99% with conservative interval coverage (95%) and strict Type-I error control (<5%) under the null. Thus, by adaptively adjusting for at most two covariates (one in the outcome regression and one in the propensity score), the TMLEs improved statistical power over the unadjusted estimator by 81%. In line with previous works,^{15,16,29} this reiterates the importance of a data-driven approach

to covariate adjustment that is tailored to maximizing efficiency. Specifically, the TMLEs differed in their selection of adjustment variables. For estimation of the outcome regression, the cluster-level TMLE selected $W1^c$ in 54% of the simulated trials and $W2^c$ in the other 46%. In contrast, the Hierarchical TMLE selected $W1$ in 82% of iterations and $W2$ in 18% of iterations. The different selections illustrate how the relationship between the cluster-level outcome and cluster-level covariates can be distinct from the relationship between the individual-level outcome and individual-level covariates—impacting the optimal adjustment strategy. Importantly, both approaches avoided adjustment for covariates that were not predictive of the outcome (ie, $\{W3^c, W4^c\}$ at the cluster-level and $\{W3, W4\}$ at the individual-level). For these estimators, additional simulation results, including fewer clusters ($J = 10$), smaller clusters (mean cluster size N_j of 20), and for the individual-level effect, are given in the Data S1. These simulations again demonstrate substantial gains in statistical power from using TMLE, while tightly preserving confidence interval coverage and Type-I error control under the null.

The other estimators, relying on fixed specifications of their outcome regressions, did improve power over their unadjusted counterparts, but not to the same extent as the TMLEs. Specifically, for estimation of the geometric incidence ratio, using CARE to adjust for both cluster-level and individual-level covariates improved statistical power to 93%, vastly surpassing the power of the t -test (14%). However, CARE failed to maintain nominal confidence interval coverage (86%). For the individual-level risk ratio, equal to the cluster-level incidence ratio in these simulations, GEE offered power improvements over the unadjusted approach (75% vs 18%, respectively). However, GEE exhibited less than nominal confidence interval coverage (89%) and inflated Type-I error rates (8%) under the null. While Aug-GEE, also targeting the individual-level risk ratio, slightly improved power over standard GEE (77%); it exhibited worse confidence interval coverage (84%) and doubled the Type-I error rate (16%).

4.3 | Simulation II

In many CRTs, participant outcomes are influenced by the cluster size. Suppose, for example, that the smallest health facilities have the fewest resources to the detriment of their patients' health, and largest health facilities are overburdened also to the detriment of their patients' health. In this setting, wide variation in cluster size can result in a divergence between cluster-level and individual-level effects. Intuitively, cluster-level parameters (eg, Equations 3 and 4) give equal weight to each cluster, regardless of its size, while individual-level parameters (eg, Equations 5 and 6) give equal weight to all trial participants. The distinction between the parameters is exacerbated when cluster size interacts with the treatment and is said to be “informative.”^{8,9} Therefore, in the second simulation study, we considered a more complex data generating process to highlight the distinction between the cluster-level effects and individual-level effects.

We again focused on a setting with $J = 20$ clusters, reflecting the PTBi study. For each cluster $j = \{1, \dots, J\}$, we generated cluster-level covariates $E2_j \sim \text{Norm}(0, 1)$, $E1_j \sim \text{Norm}(0, 1)$, and the cluster size $N_j \sim \text{Norm}(400, 250)$, again truncated at a minimum of 30 participants. For each participant $i = \{1, \dots, N_j\}$ in cluster j , we generated three individual-

level covariates $W1_{ij} \sim \text{Norm}(U_{E1_j}, 0.5)$, $W2_{ij} \sim \text{Norm}(U_{E2_j}, 0.5)$, and $W3_{ij} \sim \text{Norm}(U_{E3_j}, 0.5)$ with cluster-specific means $U_{E1_j} \sim \text{Unif}(-1, 1)$, $U_{E2_j} \sim \text{Unif}(-1, 1)$, and $U_{E3_j} \sim \text{Unif}(-1, 1)$.

As before, clusters were pair-matched on $E2$, and within each pair, one cluster was randomized to the intervention arm ($A = 1$) and the other to the control arm ($A = 0$). Lastly, we simulated the individual-level outcomes as a function of the treatment, the cluster-level and individual-level covariates, and unmeasured factor $U_{Y_{ij}} \sim \text{Unif}(0, 1)$:

$$Y_{ij} \sim \mathbb{1}[U_{Y_{ij}} < \text{logit}^{-1}(0.5 + W1_{ij}/6 + W2_{ij}/2 + W3_{ij}/4 + E1_j/5 + E2_j/5 - \tilde{N}_j/8 - A\tilde{N}_j/5)],$$

where $\tilde{N}_j = N_j / 150$ denotes the scaled cluster size. Unlike Simulation I, the probability of the individual-level outcome Y_{ij} was a function of cluster size N_j . Specifically, the outcome risk was lower for participants in larger clusters, especially large clusters in the intervention arm. As before, we generated counterfactual outcomes under the intervention and under the control by setting $A = 1$ and $A = 0$, respectively. Then for a population of 1000 clusters, we calculated the relative effect at the cluster-level (Equation 10) and at the individual-level (Equation 11).

For this simulation, we focused on the performance of the cluster-level TMLE and Hierarchical TMLE, given their flexibility to estimate a variety of effects and their ability to incorporate baseline covariates to improve precision and statistical power, while maintaining Type-I error control. As previously discussed, through the application of weights, the cluster-level TMLE can estimate individual-level effects. Likewise, Hierarchical TMLE, an estimation approach based on individual-level data, can estimate cluster-level effects. It is possible that the other CRT estimators have this flexibility, but the needed extensions remain to be fully studied.⁴⁵

In this simulation, each TMLE used Adaptive Prespecification to choose at most two covariates for adjustment if their inclusion improved efficiency as compared to an unadjusted effect estimator. The candidate adjustment set for the cluster-level TMLE included $\{\emptyset, W_1^c, W_2^c\}$, while the candidate adjustment set for individual-level TMLE included $\{\emptyset, W_1, W_2\}$. For comparison, we also considered a cluster-level TMLE with fixed adjustment for W_1^c in the cluster-level outcome regression and a Hierarchical TMLE with fixed adjustment for W_1 in the individual-level outcome regression. Inference was based on an estimate of the influence function. We again focus on the analysis breaking the matches used for randomization.

4.4 | Results of simulation II

In this simulation, the true value of the cluster-level relative effect (Equation 10) was 0.78, substantially smaller than the true value of the individual-level relative effect (Equation 11) of 0.69. In other words, the intervention resulted in a 22% relative reduction in the incidence of the outcome and a 31% relative reduction in the individual-level risk of the outcome. Of course, this is just one simulation study, and in practice, there is no guarantee that the cluster-level effect will be smaller, or even different, from the individual-level effect.

For this simulation study, Table 2 shows the performance of the cluster-level TMLE and Hierarchical TMLE with fixed and adaptive adjustment. Given the differing magnitude of the effects, it is unsurprising that estimators of the individual-level effect, shown on the right, achieved notably higher power than estimators of the cluster-level effect, shown on the left. Specifically, estimators of the cluster-level effect achieved a maximum power of 44%, while estimators of the individual-level effect achieved a maximum power of 68%.

Focusing on estimators of the cluster-level effect (Table 2, Left), all approaches resulted in low bias, similar variability (σ), and good confidence interval coverage (96%). For a given adjustment strategy (fixed or adaptive), there was little practical difference in performance between analyses using cluster-level or individual-level data. However, as expected by theory,²⁹ the TMLEs using Adaptive Prespecification achieved higher power ($\approx 44\%$) than the TMLEs relying on fixed specification of the outcome regression (40%). For the individual-level effect (Table 2, Right), the gains in power with adaptive adjustment were more substantial. Specifically, TMLEs with fixed adjustment achieved a maximum power of 60%, while the TMLEs using Adaptive Prespecification achieved a maximum power of 68%. While the approaches were slightly biased toward the null, all had good confidence interval coverage (94%-96%). Again we see little practical difference between approaches relying on cluster-level data or utilizing individual-level data. As shown in the Data S1, all approaches also maintained strong Type-I error control under the null.

Altogether the results of this simulation demonstrate the ability of TMLE to estimate both cluster-level and individual-level effects, while adaptively adjusting for baseline covariates to maximize efficiency. These results also highlight the critical importance of prespecifying the primary effect measure and using a CRT estimator of that effect. Notable bias and misleading inference may arise when the statistical estimation approach is mismatched with the desired causal effect. In all settings, the research question should drive the specification of the target effect and thereby the statistical estimation approach.^{10,12-14}

5 | REAL DATA APPLICATION: THE PTBI STUDY IN KENYA AND UGANDA

In East Africa, preterm birth remains a leading risk factor for perinatal mortality, defined as stillbirth and first-week deaths.⁵⁹ Evidence-based practices, such as use of antenatal corticosteroids and skin-to-skin contact, are not routinely used and have the potential to improve outcomes for preterm infants during the critical intrapartum and immediate newborn periods. The PTBi study was a CRT designed to improve the quality of care for mothers and preterm infants at the time of birth through a health facility intervention ([ClinicalTrials.gov: NCT03112018](https://clinicaltrials.gov/ct2/show/study/NCT03112018)).⁵⁹ The primary endpoint was intrapartum stillbirth and 28-day mortality among preterm infants delivered from October 2016 to May 2018 in Western Kenya and Eastern Uganda.

In more detail, 20 health facilities, including large hospitals and smaller health centers, were selected for participation in the study. The facilities ranged in size, staff-to-patient ratio, and capacity to perform cesarean section (C-section), among others. Prior to randomization, facilities were pair-matched on country, delivery volume, staff-to-patient ratio, as well as rates of stillbirths, low-birth weight infants, and pre-discharge neonatal mortality.⁶⁰ Within

matched pairs, they were then randomized to either the intervention or control arm.⁵⁹ Facilities in the control arm received (1) strengthening of routine data collection and (2) introduction of the WHO Safe Childbirth Checklist.⁶¹ The facilities in the intervention arm received the components included in the control arm in addition to (1) PRONTO™ Simulation training,⁶² and (2) quality improvement collaboratives aimed to reinforce and optimize use of evidence-based practices. All study components consisted of known interventions and strategies aiming to improve quality of care, teamwork, communication, and data use.⁵⁹

The results of the PTBi Study have been previously published.⁶⁰ Here, we focus on the real-world impact of highly variable cluster sizes when defining and estimating causal effect in CRTs. Specifically, during the study period, an unforeseen political strike led to lack of medical providers at certain facilities, thereby decreasing volume at some facilities while increasing volume at others. The number of preterm births for a given facility N_j ranged between 40 and 366 in the intervention arm and between 29 and 447 in the control arm. Differences in cluster size were also pronounced within matched pairs and ranged between 9 and 211.

To study the impact of high variability in cluster size, we return to the consequences of specifying the effect in terms of cluster-level outcomes (Equation 10) vs individual-level outcomes (Equation 11). In PTBi, the individual-level outcome Y_{ij} was an indicator of preterm infant mortality by 28-day follow-up. Infants dying before discharge (stillbirth and predischarge mortality) were also included in the study; for these infants, $Y_{ij} = 1$. For facility $j = \{1, \dots, 20\}$, the cluster-level endpoint Y_j^c was the incidence of fresh stillbirth and 28-day all-cause mortality among preterm births and calculated as the empirical mean of the individual-level outcomes: $Y_j^c = 1 / N_j \sum_{i=1}^{N_j} Y_{ij}$.

For both the cluster-level and individual-level effects, we compared estimates and inference from the TMLEs using cluster-level or individual-level data. The cluster-level TMLE used Adaptive Prespecification to select between no adjustment and adjustment for the proportion of mothers receiving a C-section, while the Hierarchical TMLE used Adaptive Prespecification to select between no adjustment and adjustment for an individual-level indicator of receiving a C-section. For comparison, we also implemented each TMLE without adjustment and each approach breaking or preserving the matched pairs used for randomization. (See Balzer et al.³⁵ for details on how matched analyses impact estimation and inference with TMLE.) Since C-section status was missing for 47 participants, all analyses restricted to the 2891 mother-infant dyads with complete data to improve comparability of the methods in this demonstration paper.

5.1 | PTBi results

As shown in Table 3, estimates and inference for the effect of the PTBi intervention varied substantially by the target of inference. At the cluster-level, the average incidence of 28-day mortality among preterm infants was 12% among facilities randomized to the intervention and 15% among facilities randomized to the control. Therefore, the unadjusted estimator of the cluster-level relative effect (Equation 10) was 0.81, corresponding to 19% reduction

in the incidence of mortality at the facility-level. The results for the individual-level effect were markedly different, reflecting how outcome risk varied by cluster size. As shown in Figure 1, larger hospitals tended to have poorer outcomes (Pearson's correlation $r=0.77$), especially in the control arm ($r=0.88$). At an individual-level, the overall proportion of preterm infants who died within 28 days was 15% in the intervention arm and 23% in the control arm. Therefore, the unadjusted estimator of the individual-level relative effect (Equation 11) was 0.66, corresponding to a 34% reduction in the mortality risk. As predicted by theory,^{1,33,35} for a given statistical estimand, the analysis preserving the matched pairs used for randomization was notably more precise than the analysis breaking the matches. Specifically, when keeping vs breaking the matches, the unadjusted estimator was five times more efficient for the cluster-level effect and three times more efficient for the individual-level effect. Throughout, efficiency is defined as the variance of the unadjusted effect estimator when breaking the matches divided by the variance of another approach.¹⁶

Estimates from the TMLEs also indicated that the PTBi intervention reduced preterm mortality; again, the results varied by the target of inference and the estimation approach. For the cluster-level effect (Table 3, Left), adaptive adjustment for C-section doubled the precision of the analyses when breaking the matches. However, when preserving the matches, both the cluster-level TMLE and the Hierarchical TMLE defaulted to the unadjusted estimator, reflecting the ability of the adaptive approach to adjust only when it improves efficiency. Here, further adjustment for C-section status did not improve efficiency after matching on key outcome predictors (eg, region and outcome rates prior to randomization). Similar results are seen for the intervention effect at the individual-level effect (Table 3, Right). Adjustment for C-section improved the precision of analyses breaking the matches, but not analyses preserving the matches used for randomization. It is worth noting that adjusted estimates tended to be of smaller magnitude than the unadjusted ones, especially for the cluster-level TMLE. This is not always the case,¹⁶ and we return to this in detail below.

6 | DISCUSSION

CRTs are commonly used to evaluate the effects of interventions, which are randomized to groups of individuals, such as communities, clinics, or schools. In all trials, it is essential to clearly define the target causal effect and ensure the estimation approach follows from that effect.¹⁴ However, in CRTs, the hierarchical data structure (eg, participants nested within clinics) complicates effect specification, especially when cluster sizes vary and interact with the intervention effect. Additionally, in both resource-rich and resource-limited settings, high costs and implementation barriers often limit the number of clusters that can be enrolled and randomized. Therefore, it is important to understand the potential advantages and pitfalls of common analytic approaches in CRTs, overall and when there are few independent units of varying size.

In this article, we aimed to provide a comprehensive framework for defining causal effects in CRTs and an in-depth comparison of methods to estimate those effects. First, we used a structural causal model to describe the data generating process of a CRT and intervened upon it to generate counterfactual outcomes. Using these counterfactuals, we explored a

variety of causal parameters that could be of interest, and we stressed the importance of a priori-specifying the primary effect measure. We then delved into common CRT methods, and considered each theoretically and with finite sample simulations. Finally, we demonstrated the impact of key analytic decisions using real data from the PTBi study, a clinic-based CRT designed to reduce mortality among preterm infants in East Africa.

Our theoretical comparison of estimators revealed that common methods often target different causal effects. Results from our simulation studies demonstrated how TMLE could improve statistical power by adaptively adjusting for covariates, while maintaining nominal confidence interval coverage. In the first simulation study, approaches based on generalized estimating equations (GEE and Aug-GEE) failed to preserve Type-I error control; this is consistent with the literature warning against their application with fewer than 30 clusters.^{1,7} However, it is worth noting that some studies may value minimizing Type-II error over Type-I error, and this should be carefully considered when developing the statistical analysis plan. Results from our second simulation study demonstrated the impact of informative cluster size, which occurs when the cluster size interacts with the intervention effect. In this setting, there can be a sharp divergence in the interpretation and magnitude of effects defined at the cluster-level or individual-level. This divergence was observed in the real data application, where larger facilities in the control arm had poorer outcomes (Figure 1).

There are several limitations to the applied analysis, which merit additional consideration. Primarily, we only addressed two methodological challenges in PTBi: defining and estimating causal effects in CRTs with few clusters of varying sizes. We did not address how to select candidate adjustment variables in practice.^{15,29} Instead, due to limited data availability, our investigation explored adjustment for a single covariate: having a C-section at the individual-level or the proportion of women receiving a C-section at the cluster-level. This is a complex covariate in PTBi for the following reasons.⁶⁰ First, not all health facilities had C-section capabilities at baseline. Further, the capacity to perform C-section changed for some facilities during the trial. Second, women with more complicated pregnancies were referred to facilities with C-section capacity, regardless of the randomization arm. However, during the trial, a political strike occurred, substantially impacting where women delivered—again regardless of randomization arm. These real-world complexities reflect how CRTs are pragmatic and often need to respond to exogenous factors influencing both study interventions and outcomes. These secular events may entail modifications to the study design, follow-up period, primary endpoint, and statistical analysis plan; see, for example, Kakande et al.⁶³

As previously discussed, adjustment for C-section in PTBi attenuated estimates of the intervention effect, especially at the cluster-level. It is possible that the intervention impacted who and where women received C-sections—thus violating the exclusion restriction in the causal model (Equation 1). While formal significance testing indicated no such effect, we cannot rule out that C-section might actually be a mediator. This reflects a wider challenge, common in CRTs: how to define “baseline” when a longitudinal cohort of participants is not formed immediately after cluster-level randomization. Instead, in CRTs, participants may enroll or have their outcomes measured at various times after cluster-level randomization. Specifically in PTBi, mothers enrolled at the time of delivery, and infant outcomes were then

measured from delivery until 28-day post-delivery. Addressing challenges with potential mediators, time-varying covariates, and missingness is beyond the scope of this review paper and can be handled using more complex methods.¹⁶⁻²³

There are also limitations in our methods comparison. First, while we provided a comprehensive framework for defining causal effects, we did not provide explicit recommendations on whether to target the cluster-level or individual-level effect, on the additive or relative scale, and for the study units or a wider population. Instead, specification of the research query and the causal effect of interest must be made as a study team and will vary across real data applications. As we have demonstrated, TMLE provides flexibility to estimate a wide variety of user-specified effects in CRTs. In simulations, there was little practical difference in the performance of the cluster-level TMLE, using cluster-level data, vs Hierarchical TMLE, using individual-level data, for the same statistical estimand. Both incorporated Adaptive Prespecification to select the optimal adjustment approach, given the data structure and target of inference. However, in the PTBi study, there were substantive differences in the point estimates and inference, when using the cluster-level TMLE vs Hierarchical TMLE. This may reflect that while both cluster-level and Hierarchical TMLEs can target the same parameter, the relationship between the cluster-level covariates and outcome may be distinct from the relationship between the individual-level covariates and outcome. Specifically, the impact of the proportion of women receiving a C-section on the cluster-level incidence of mortality is likely to be different than the impact of an individual woman's having a C-section on the risk of mortality for her child. Future work will extend Adaptive Prespecification to select between cluster-level and Hierarchical TMLEs, whichever maximizes empirical efficiency for the effect of interest. For now, the choice between cluster-level TMLE and Hierarchical TMLE (for the same parameter) may be driven by data availability and ease of implementation. Practically, it is most straightforward to estimate cluster-level effects with a cluster-level TMLE and individual-level effects with Hierarchical TMLE. However, in certain applications, data may only be available at the cluster-level, necessitating the use of cluster-level TMLE with weights to estimate individual-level effects. Altogether, to help inform the optimal analysis for a given application, we recommend conducting a simulation study reflecting the nuances of the real data problem.

In summary, CRTs are a popular approach for researchers seeking to estimate effects when interventions are scaled up to the population level. Many of the challenges faced by the PTBi study are common and highlight the importance of carefully selecting and prespecifying the primary causal effect, which should be driven by the trial's objectives.¹⁴ For statistical estimation and inference, we recommend using TMLE given its flexibility to estimate a variety of effects and data-adaptively adjust for both cluster-level and individual-level covariates, while preserving Type-I error control. We hope our presentation and comparison of various TMLEs has helped clarify its use and benefits in a CRT setting. Additionally, we have provided full R code in the Data S1 and note that a R package for more general use of TMLE in randomized trials is under way.⁶⁴

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

The authors thank the women and infants who participated in the PTBi study. They thank the community health volunteers and the providers at each study facility in Kenya and Uganda, the PTBi research teams at UCSF, Makerere University, and Kenya Medical Research Institute. Authors also thank Gertrude Namazzi, Kevin Achola, Christopher Otare, Paul Mubiri, Rikita Merai, Nancy Sloan, Lara Miller, Wenjing Zheng, and Caleb Miles, the Bill and Melinda Gates Foundation and National Institutes of Health for their generous support of this work. Lastly, the authors are very grateful for the mentors at UC Berkeley School of Public Health.

FUNDING INFORMATION

This work was supported by the National Institutes of Health (U01AI150510, UM1AI068636, R01AI125000, R01AI074345, 2R01AI074345-10A1). This work was also supported by the Bill & Melinda Gates Foundation (OPP1107312). Under the grant conditions of the Foundation, a Creative Commons Attribution 4.0 Generic License has already been assigned to the Author Accepted Manuscript version that might arise from this submission. The funders did not play any direct role in study design, data collection, analysis, or writing of the manuscript.

DATA AVAILABILITY STATEMENT

Computing code to reproduce the simulation studies and to analyze the PTBi Study is available at https://github.com/LauraBalzer/Comparing_CRT_Methods.

REFERENCES

- Hayes RJ, Moulton LH. Cluster Randomised Trials. 2nd ed. Boca Raton, FL, USA: Chapman & Hall/CRC; 2017.
- Turner EL, Li F, Gallis JA, Prague M, Murray DM. Review of recent methodological developments in group-randomized trials: part 1—design. *Am J Public Health*. 2017;107:907–915. [PubMed: 28426295]
- Turner EL, Prague M, Gallis JA, Li F, Murray DM. Review of recent methodological developments in group-randomized trials: part 2—analysis. *Am J Public Health*. 2017;107:7. [PubMed: 27925807]
- Crespi CM. Improved designs for cluster randomized trials. *Annu Rev Public Health*. 2016;37(1):1–16.
- Murray DM, Taljaard M, Turner EL, George SM. Essential ingredients and innovations in the design and analysis of group-randomized trials. *Annu Rev Public Health*. 2020;41:1–19. [PubMed: 31869281]
- Pagel C, Prost A, Lewycka S, et al. Intracluster correlation coefficients and coefficients of variation for perinatal outcomes from five cluster-randomised controlled trials in low and middle-income countries: results and methodological implications. *Trials*. 2011;12:151. [PubMed: 21672223]
- Murray DM, Pals SL, George SM, et al. Design and analysis of group-randomized trials in cancer—a review of current practices. *Prev Med*. 2018;111:241–247. [PubMed: 29551717]
- Seaman SR, Pavlou M, Copas AJ. Review of methods for handling confounding by cluster and informative cluster size in clustered data. *Stat Med*. 2014;33:5371–5387. [PubMed: 25087978]
- Nevalainen J, Datta Oja H. Inference on the marginal distribution of clustered data with informative cluster size. *Stat Papers*. 2014;55(1):71–92.
- Petersen ML, Van der Laan MJ. Causal models and learning from data: integrating causal modeling and statistical estimation. *Epidemiology*. 2014;25(3):418–426. [PubMed: 24713881]
- van der Laan MJ, Rose S. Targeted Learning: Causal Inference for Observational and Experimental Data. 1st ed. Berlin, Germany: Springer Science+Business Media; 2011.
- Balzer LB, Petersen ML. Machine learning in causal inference: how do I love thee? Let me count the ways. *Am J Epidemiol*. 2021;190(8):1483–1487. [PubMed: 33751059]

13. Kahan BC, Morris TP, White IR, Carpenter J, Cro S. Estimands in published protocols of randomised trials: urgent improvement needed. *Trials*. 2021;22(686):686. 10.1186/s13063-021-05644-4 [PubMed: 34627347]
14. International Council for Harmonisation. ICH E9 (R1) Addendum on Estimands and Sensitivity Analysis in Clinical Trials to the Guideline on Statistical Principles for Clinical Trials. Amsterdam, Netherlands: European Medicines Agency; 2020. <https://www.ema.europa.eu/en/ich-e9-statistical-principles-clinical-trials>
15. Benkeser D, Díaz I, Luedtke A, Segal J, Scharfstein D, Rosenblum M. Improving precision and power in randomized trials for COVID-19 treatments using covariate adjustment, for binary, ordinal, and time-to-event outcomes. *Biometrics*. 2021;77(4):1467–1481. [PubMed: 32978962]
16. Balzer LB, Van der Laan M, Ayieko J, et al. Two-stage TMLE to reduce bias and improve efficiency in cluster randomized trials. *Biostatistics*. 2023;24:502–517. [PubMed: 34939083]
17. Díaz-Ordaz K, Kenward MG, Cohen A, Coleman CL, Eldridge S. Are missing data adequately handled in cluster randomised trials? A systematic review and guidelines. *Clin Trials*. 2014;11(5):590–600. [PubMed: 24902924]
18. Fiero MH, Huang S, Oren E, Bell ML. Statistical analysis and handling of missing data in cluster randomized trials: a systematic review. *Trials*. 2016;17:72. [PubMed: 26862034]
19. Balzer LB, Ayieko J, Kwarisiima D, et al. Far from MCAR: obtaining population-level estimates of HIV viral suppression. *Epidemiology*. 2020;31(5):620–627. [PubMed: 32452912]
20. Leyrat C, Caille A, Donner A, Giraudeau B. Propensity scores used for analysis of cluster randomized trials with selection bias: a simulation study. *Stat Med*. 2013;32(19):3357–3372. [PubMed: 23553813]
21. Leyrat C, Caille A, Donner A, Giraudeau B. Propensity score methods for estimating relative risks in cluster randomized trials with low incidence binary outcomes and selection bias. *Stat Med*. 2014;33(20):3556–3575. [PubMed: 24771662]
22. Li F, Tian Z, Bobb J, Papadogeorgou G, Li F. Clarifying selection bias in cluster randomized trials. *Clin Trials*. 2022;19(1):33–41. [PubMed: 34894795]
23. Nugent JR, Marquez C, Charlebois ED, Abbott R, Balzer LB. Blurring cluster randomized trials and observational studies: two-Stage TMLE for sub-sampling, missingness, and few independent units. 2022. <https://arxiv.org/abs/2208.09508>
24. Li F, Lokhnygina Y, Murray DM, Heagerty PJ, DeLong ER. An evaluation of constrained randomization for the design and analysis of group randomized trials. *Stat Med*. 2016;35(10):1565–1579. [PubMed: 26598212]
25. Li F, Turner EL, Heagerty PJ, Murray DM, Vollmer WM, DeLong ER. An evaluation of constrained randomization for the design and analysis of group randomized trials with binary outcomes. *Stat Med*. 2017;36(24):3791–3806. [PubMed: 28786223]
26. Liang K-Y, Zeger SL. Longitudinal data analysis using generalized linear model. *Biometrika*. 1986;73(1):13–22.
27. Gail MH, Mark SD, Carroll RJ, Green SB, Pee D. On design considerations and randomization-based inference for community intervention trials. *Stat Med*. 1996;15:1069–1092. [PubMed: 8804140]
28. Balzer L, Zheng W, Van der Laan MJ, Petersen ML. A new approach to hierarchical data analysis: targeted maximum likelihood estimation for the causal effect of a cluster-level exposure. *Stat Methods Med Res*. 2019;28(6):1761–1780. [PubMed: 29921160]
29. Balzer L, van der Laan MJ, Petersen M. Adaptive pre-specification in randomized trials with and without pair-matching. *Stat Med*. 2016;35:4528–4545. [PubMed: 27436797]
30. Stephens AJ, Tchetgen Tchetgen EJ, de Gruttola V. Augmented GEE for improving efficiency and validity of estimation in cluster randomized trials by leveraging cluster and individual-level covariates. *Stat Med*. 2012;31(10):915–930. [PubMed: 22359361]
31. Bellamy SL, Gibberd R, Hancock L, et al. Improved designs for cluster randomized trials. *Stat Methods Med Res*. 2000;9:135–159. [PubMed: 10946431]
32. Su F, Ding P. Model-assisted analyses of cluster-randomized experiments. *J R Stat Soc Series B Stat Methodology*. 2021;83(5):994–1015.

33. Imai K, King G, Nall C. The essential role of pair matching in cluster-randomized experiments, with application to the Mexican universal health insurance evaluation. *Stat Sci*. 2009;24(1):29–53.
34. Cook AJ, Delong E, Murray DM, Vollmer WM, Heagerty PJ. Statistical lessons learned for designing cluster randomized pragmatic clinical trials from the NIH health care systems collaboratory biostatistics and design core. *Clin Trials*. 2016;13(5):504–512. [PubMed: 27179253]
35. Balzer LB, Petersen ML, van der Laan MJ, The SEARCH Consortium. Adaptive pair-matching in randomized trials with unbiased and efficient effect estimation. *Stat Med*. 2015;34(6):999–1011. [PubMed: 25421503]
36. Pearl J. *Causality: Models, Reasoning and Inference*. 2nd ed. Cambridge, United Kingdom: Cambridge University Press; 2009.
37. Imbens GW. Nonparametric estimation of average treatment effects under exogeneity: a review. *Rev Econ Stat*. 2004;86(1):4–29. doi:10.1162/003465304323023651
38. Imai K. Variance identification and efficiency analysis in randomized experiments under the matched-pair design. *Stat Med*. 2008;27:4857–4873. [PubMed: 18618425]
39. Neyman J. Sur les applications de la theorie des probabilites aux experiences agricoles: essai des principes (In polish). English translation by Dabrowska DM and Speed TP (1990). *Stat Sci*. 1923;5:463–472.
40. Rubin DB. Comment: Neyman (1923) and causal inference in experiments and observational studies. *Stat Sci*. 1990;5(4):472–480.
41. Balzer LB, Petersen ML, van der Laan M, SEARCH Collaboration. Targeted estimation and inference for the sample average treatment effect in trials with and without pair-matching. *Stat Med*. 2016;35:3717–3732. [PubMed: 27087478]
42. Rothman KJ, Greenland S, Lash TL. *Modern Epidemiology*. 3rd. ed. Philadelphia, PA, USA: Lippincott Williams & Wilkins; 2008.
43. Balzer LB. “All generalizations are dangerous, even this one.” - Alexandre dumas [commentary]. *Epidemiology*. 2017;28(4):562–566. [PubMed: 28346266]
44. Rubin DB, van der Laan MJ. Empirical efficiency maximization: improved locally efficient covariate adjustment in randomized experiments and survival analysis. *Int J Biostat*. 2008;4(1):5. [PubMed: 19381345]
45. Wang X, Turner EL, Li F, et al. Two weights make a wrong: cluster randomized trials with variable cluster sizes and heterogeneous treatment effects. *Contemp Clin Trials*. 2022;114:106702. [PubMed: 35123029]
46. Moore KL, van der Laan MJ. Covariate adjustment in randomized trials with binary outcomes: targeted maximum likelihood estimation. *Stat Med*. 2009;28(1):39–64. doi:10.1002/sim.3445 [PubMed: 18985634]
47. Rosenblum M, van der Laan M. Simple, efficient estimators of treatment effects in randomized trials using generalized linear models to leverage baseline variables. *Int J Biostat*. 2010;6(1):13.
48. Havlir DV, Balzer LB, Charlebois ED, et al. HIV testing and treatment with the use of a community health approach in rural Africa. *N Engl J Med*. 2019;381:219–229. [PubMed: 31314966]
49. Balzer LB, Cai C, Garraza LG, Amaranath P. Adaptive selection of the optimal strategy to improve precision and power in randomized trials 2022. <https://arxiv.org/abs/2210.17453>
50. Fitzmaurice GM, Laird NM, Ware JH. *Applied Longitudinal Analysis*. 2nd ed. Hoboken, NJ, USA: Wiley; 2011.
51. Schnitzer ME, van der Laan MJ, Moodie EEM, Platt RW. Effect of breastfeeding on gastrointestinal infection in infants: a targeted maximum likelihood approach for clustered longitudinal data. *Ann Appl Stat*. 2014;8(2):703–725. [PubMed: 25505499]
52. Hubbard AE, Ahern J, Fleischer NL, et al. To GEE or not to GEE: comparing population average and mixed models for estimating the associations between neighborhood risk factors and health. *Epidemiology*. 2010;21(4):467–474. [PubMed: 20220526]
53. Li P, Redden DT. Small sample performance of bias-corrected sandwich estimators for cluster-randomized trials with binary outcomes. *Stat Med*. 2015;34(2):281–296. [PubMed: 25345738]
54. Fay MP, Graubard BI. Small-sample adjustments for Wald-type tests using sandwich estimators. *Biometrics*. 2001;57:1198–1206. [PubMed: 11764261]

55. Stephens AJ, Tchetgen Tchetgen EJ, de Gruttola V. Locally efficient estimation of marginal treatment effects when outcomes are correlated: is the prize worth the chase? *Int J Biostat*. 2014;10(1):59–75. [PubMed: 24566369]
56. Beck C, Lu B, Greevy R. nbpMatching: functions for optimal non-bipartite matching. R package version 1.5.1 2016. <https://CRAN.R-project.org/package=nbpMatching>. Accessed March 10, 2022.
57. Wang M. geesmv: modified variance estimators for generalized estimating equations. R package version 1.3. 2015. <https://CRAN.R-project.org/package=geesmv>. Accessed March 10, 2022.
58. Prague M. CRTgeeDR: doubly robust inverse probability weighted augmented GEE estimator. R package version 2.0 2017. <https://CRAN.R-project.org/package=CRTgeeDR>. Accessed March 10, 2022.
59. Otieno P et al. Strengthening intrapartum and immediate newborn care to reduce morbidity and mortality of preterm infants born in health facilities in Migor County, Kenya and Busoga region, Uganda: a study protocol for a randomized control trial. *Annu Rev Public Health*. 2018;19:313.
60. Walker D, Otieno P, Butrick E, et al. Preterm birth initiative Kenya and Uganda implementation research collaborative. Effect of a quality improvement package for intrapartum and immediate newborn care on fresh stillbirth and neonatal mortality among preterm and low-birthweight babies in Kenya and Uganda: a cluster-randomised facility-based trial. *Lancet Glob Health*. 2020;8(8):e1061–e1070. [PubMed: 32710862]
61. Spector JM, Lashoher A, Agrawal P, et al. Designing the WHO safe childbirth checklist program to improve quality of care at childbirth. *Int J Gynecol Obstet*. 2013;122(2):164–168.
62. Walker D, Cohen SR, Fritz J, et al. Impact evaluation of PRONTO Mexico: a simulation-based program in obstetric and neonatal emergencies and team training. *Simul Healthc*. 2016;11(1):1–9. [PubMed: 26312613]
63. Kakande E, Christian C, Balzer LB, et al. A mid-level health manager intervention to promote uptake of isoniazid preventive therapy in Uganda: a cluster randomized trial. *Lancet HIV*. 2022;9(9):E607–E616. [PubMed: 35908553]
64. Balzer LB, Nugent J. tml4CRTs: TMLE for randomized trials. 2022.

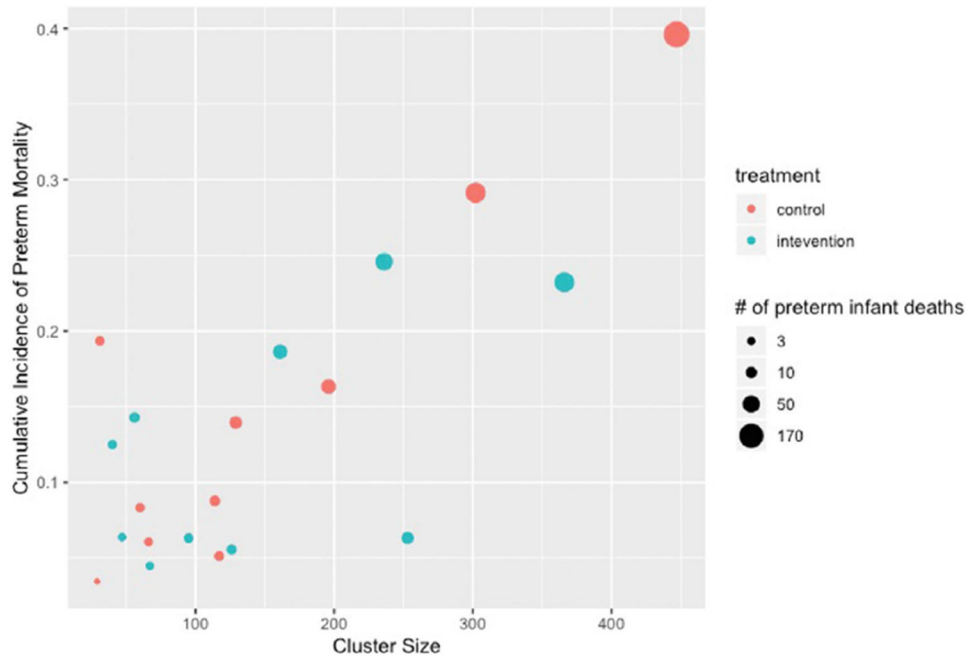


FIGURE 1. Scatter plot of cluster size N_j by the cumulative incidence of preterm mortality and trial arm in Preterm Birth Initiative.

TABLE 1

Performance of common cluster randomized trial estimators when there is an effect and under the null across 500 iterations of Simulation I.

	When there is an effect					Under the null						
	pt	bias	σ	$\bar{\sigma}$	covg	power	pt	bias	σ	$\bar{\sigma}$	covg	Type-I
Unadj	0.84	0.01	0.17	0.17	0.97	0.18	1.01	0.01	0.16	0.16	0.96	0.04
C-TMLE	0.83	0.00	0.04	0.05	0.98	0.99	1.00	0.00	0.03	0.03	0.98	0.02
H-TMLE	0.83	0.00	0.04	0.05	0.98	0.99	1.00	0.00	0.03	0.04	0.98	0.02
t-test	0.83	0.02	0.22	0.22	0.96	0.14	1.02	0.02	0.21	0.21	0.96	0.04
CARE	0.83	0.02	0.05	0.05	0.86	0.93	1.00	0.00	0.03	0.03	0.96	0.04
GEE	0.82	-0.01	0.08	0.07	0.89	0.75	1.00	0.00	0.08	0.07	0.92	0.08
A-GEE	0.83	-0.00	0.09	0.06	0.84	0.77	1.00	0.00	0.09	0.06	0.84	0.16

Abbreviations: "A-GEE," Augmented-GEE; bias, the average difference between the point estimate and the target effect; "covg", the proportion of times the 95% confidence interval contained the true effect; "C-TMLE" and "H-TMLE", the cluster-level TMLE and Hierarchical TMLE, respectively; both were implemented with Adaptive Prespecification; "pt", the average point estimate; "power", the proportion of times the false null hypothesis was rejected; σ , the SD of the point estimates on the log-scale; $\bar{\sigma}$, the average SE estimate on the log-scale; "Type-I", the proportion of times the true null hypothesis was rejected; "Unadj", the unadjusted estimator, implemented as a cluster-level targeted maximum likelihood estimation (TMLE) for the effect of interest and without covariate adjustment.

TABLE 2

Performance of targeted maximum likelihood estimations (TMLEs) for the cluster-level and individual-level relative effects across 500 iterations in Simulation II.

	Cluster-level effect: $\Psi_0^c(1) / \Psi_0^c(0) = 0.78$						Individual-level effect: $\Psi_0(1) / \Psi_0(0) = 0.69$					
	pt	bias	σ	$\hat{\sigma}$	covg	power	pt	bias	σ	$\hat{\sigma}$	covg	power
C-TMLE	0.78	0.00	0.13	0.14	0.97	0.40	0.71	0.02	0.15	0.15	0.95	0.59
C-TMLE-AP	0.79	0.01	0.12	0.12	0.96	0.44	0.71	0.02	0.14	0.14	0.94	0.68
H-TMLE	0.78	0.00	0.12	0.14	0.98	0.40	0.71	0.02	0.15	0.16	0.96	0.60
H-TMLE-AP	0.79	0.01	0.11	0.12	0.97	0.43	0.72	0.03	0.13	0.14	0.95	0.65

Abbreviations: “bias”, the average difference between the point estimate and the target effect; “covg”, the proportion of times the 95% confidence interval contained the true effect; “C-TMLE” and “H-TMLE”, the cluster-level TMLE and Hierarchical TMLE, respectively. Both were implemented with fixed or adaptive adjustment via Adaptive Prespecification (“-AP”); “power”, the proportion of times the false null hypothesis was rejected; “pt”, the average point estimate; σ , the SD of the point estimates on the log-scale; $\hat{\sigma}$, the average SE estimate on the log-scale.

TABLE 3

Estimating cluster-level and individual-level relative effects in the Preterm Birth Initiative (PTBi) study.

	For the cluster-level effect					For the individual-level effect				
	$\widehat{\Psi}^c(1)$	$\widehat{\Psi}^c(0)$	Ratio (95% CI)	Eff.	Adj.	$\widehat{\Psi}(1)$	$\widehat{\Psi}(0)$	Ratio (95% CI)	Eff.	Adj.
Breaking the matches										
Unadj.	12%	15%	0.81 (0.43-1.55)	1	—	15%	23%	0.66 (0.33-1.31)	1	—
C-TMLE	13%	14%	0.88 (0.59-1.34)	2	C-sect	17%	21%	0.84 (0.53-1.34)	2	C-sect
H-TMLE	12%	15%	0.84 (0.50-1.41)	2	C-sect	16%	22%	0.70 (0.38-1.27)	1	C-sect
Preserving the matches										
Unadj.	12%	15%	0.81 (0.59-1.11)	5	—	15%	23%	0.66 (0.44-0.99)	3	—
C-TMLE	12%	15%	0.81 (0.59-1.11)	5	∅	15%	23%	0.65 (0.43-0.98)	3	∅
H-TMLE	12%	15%	0.82 (0.60-1.11)	5	∅	16%	22%	0.70 (0.46-1.06)	3	C-sect

Abbreviations: “C-TMLE” and “H-TMLE”, the cluster-level TMLE and Hierarchical TMLE, respectively; both were implemented with Adaptive Prespecification to adjust for C-section (“C-sect”) or nothing (∅); “Eff”, the relative efficiency; the variance estimate for the unadjusted effect estimator breaking the matches used for randomization, divided by the variance estimate of another approach (eg, Hierarchical-TMLE with Adaptive Prespecification, keeping the matches used for randomization); “Unadj.”, the unadjusted effect estimator.