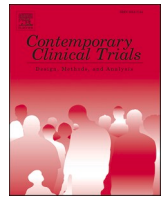




Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Elsevier hereby grants permission to make all its COVID-19-related research that is available on the COVID-19 resource centre - including this research content - immediately available in PubMed Central and other publicly funded repositories, such as the WHO COVID database with rights for unrestricted research re-use and analyses in any form or by any means with acknowledgement of the original source. These permissions are granted for free by Elsevier for as long as the COVID-19 resource centre remains active.



Causal inference methods for small non-randomized studies: Methods and an application in COVID-19

Sarah Friedrich^{*}, Tim Friede

Department of Medical Statistics, University Medical Center Göttingen, Humboldtallee 32, 37073 Göttingen, Germany

ARTICLE INFO

Keywords:
 COVID-19
 Causal inference
 Propensity score
 Small samples

ABSTRACT

The usual development cycles are too slow for the development of vaccines, diagnostics and treatments in pandemics such as the ongoing SARS-CoV-2 pandemic. Given the pressure in such a situation, there is a risk that findings of early clinical trials are overinterpreted despite their limitations in terms of size and design. Motivated by a non-randomized open-label study investigating the efficacy of hydroxychloroquine in patients with COVID-19, we describe in a unified fashion various alternative approaches to the analysis of non-randomized studies. A widely used tool to reduce the impact of treatment-selection bias are so-called propensity score (PS) methods. Conditioning on the propensity score allows one to replicate the design of a randomized controlled trial, conditional on observed covariates. Extensions include the g-computation approach, which is less frequently applied, in particular in clinical studies. Moreover, doubly robust estimators provide additional advantages. Here, we investigate the properties of propensity score based methods including three variations of doubly robust estimators in small sample settings, typical for early trials, in a simulation study. R code for the simulations is provided.

1. Introduction

Pandemic situations such as the currently ongoing SARS-CoV-2 pandemic require the fast development of diagnostics, vaccines and treatments. As the usual development programs are too long in these situations, more efficient development pathways are sought. These include more innovative approaches such as platform trials and adaptive designs [64]. Furthermore, in situations of desperate medical need such as with COVID-19, early clinical trials might receive more attention than they would normally do. Ferreira et al. call this a “disruption of medical and scientific paradigms” [31]. In March 2020, for instance, Gautret et al. [36] published a report of a small open-label non-randomized controlled study suggesting that “hydroxychloroquine treatment is significantly associated with viral load reduction/disappearance in COVID-19 patients”. Although typically not much notice would have been taken of such a small-scale study with its methodological limitations, the treatment was hailed “a game changer” by the US president putting pressure on the regulatory authorities to license the drug for COVID-19 [55].

In particular when a lot of importance is placed on non-randomized studies, their analyses and interpretation must be robust. Non-

randomized studies might be prone to bias due to confounding. One common approach to deal with this is covariate adjustment in regression models. In clinical trial applications with a binary outcome, logistic regression is usually the method of choice. However, in the case of small sample sizes, the number of possible variables to adjust for is limited by the observed events. Moreover, the use of odds ratios is not without criticism in the literature [19,25,27,63] and it is often argued that the risk difference is of greater importance to clinical decision makers [12].

Besides covariate adjustment a wide range of methods were proposed to deal with confounding. A widely used tool to reduce the impact of treatment-selection bias in observational data are so-called propensity score (PS) methods. The propensity score is defined as a participant’s probability of receiving treatment given the observed covariates [56,57]. Conditioning on the propensity score allows one to replicate the design of a randomized controlled trial, conditional on *observed* covariates. Extensions include the g-computation [38,53], which is less frequently applied, in particular in clinical studies. Moreover, doubly robust estimators have been proposed. Here, it is sufficient that either the outcome or the propensity score model is correctly specified. Hence, a doubly robust estimator does not rely on correct specification of both models.

^{*} Corresponding author.

E-mail addresses: sarah.friedrich@med.uni-goettingen.de (S. Friedrich), tim.friede@med.uni-goettingen.de (T. Friede).

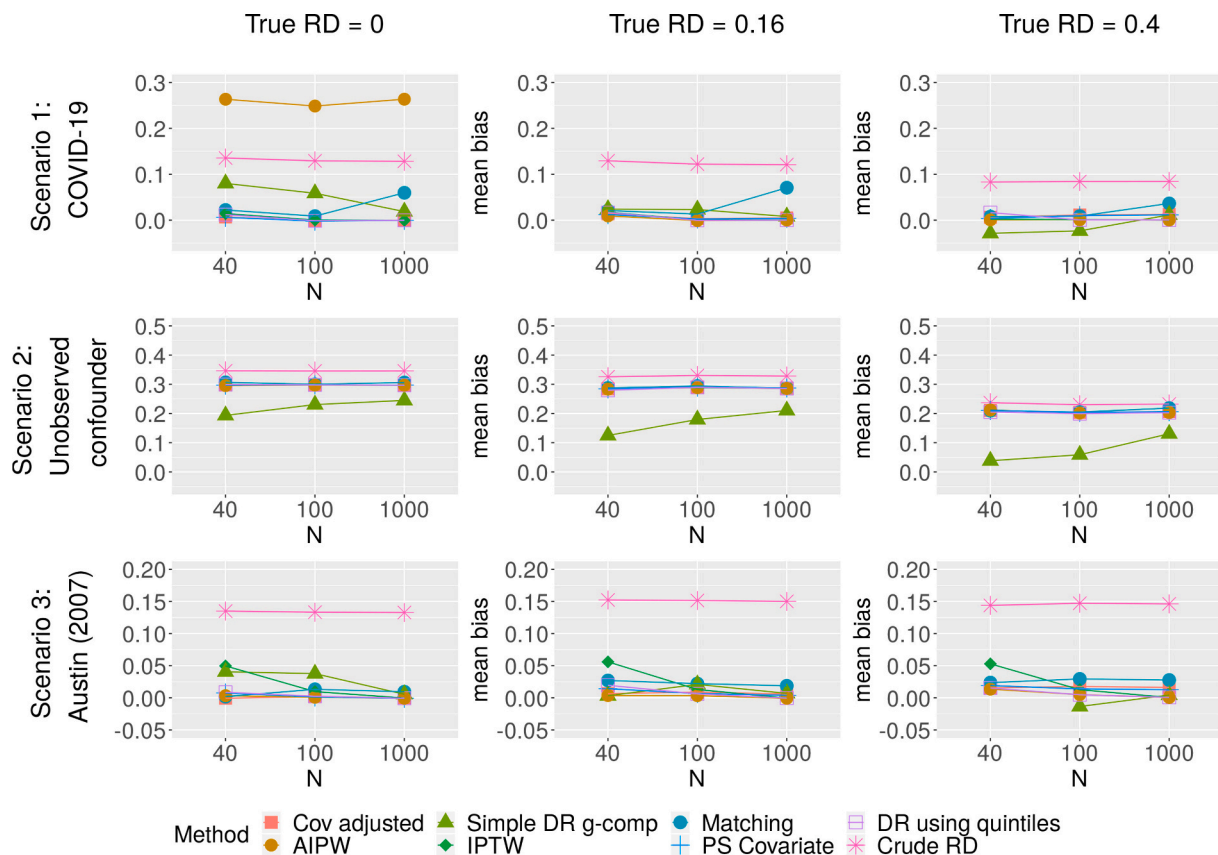


Fig. 1. Displayed is the mean bias for the three scenarios (rows) and the three simulated risk differences (columns).

Gautret et al. [36] did not apply any of these methods for non-randomized studies, but analyzed the trial as if it was randomized. Here, we describe in a unified fashion various alternative approaches and explore in simulations whether different methods might have led to different conclusions. New evidence has emerged in the meanwhile and we now know that hydroxychloroquine is not an appropriate therapy in COVID-19 [23,60]. Thus, we wonder whether a more appropriate analysis of the study by Gautret et al. [36] could have prevented much of the hype and as a result might have saved some resources.

In the context of the analysis of clinical registries and routine data including electronic health records some of the methods described above have widely been applied and their characteristics explored in simulation studies. Given the applications, simulation experiments naturally considered large-scale data sets [14]. To the best of our knowledge, propensity score methods for small samples have received less attention in the literature so far. Recently, Pirracchio et al. [51] compared PS matching and weighting estimators in small sample populations and Andrillon et al. [9] investigated properties of different matching algorithms. Here, we investigate the properties of propensity score based methods including g-computation in small sample settings, typical for early trials, in a simulation study. The manuscript is organized as follows. In Section 2 we introduce the study by Gautret et al. [36], which motivated our investigations, in more detail. In Section 3 several approaches to the analysis of non-randomized trials are described. Their properties are assessed in a simulation study, in particular in the setting of small sample sizes, in Section 4. We close with a brief discussion of the findings and the limitations of our study (Section 5).

2. Motivating example in COVID-19

Gautret et al. [36] conducted an open-label non-randomized study

investigating the efficacy and safety of hydroxychloroquine in addition to standard of care in comparison to standard of care alone. The patients in the hydroxychloroquine group were all from the coordinating centre whereas the controls were recruited from several centres including the coordinating centre. In the coordinating centre, however, only those patients refusing therapy with hydroxychloroquine were included as controls. A total of 36 patients were included in the analyses, 20 patients receiving hydroxychloroquine and 16 control patients. Out of the 20 patients on hydroxychloroquine, 6 patients received in addition also azithromycin. For the purpose of illustration, we only consider two treatment groups, i.e. with and without hydroxychloroquine. The primary outcome was virological clearance at Day 6 (with Day 0 being baseline). The individual participant data of the study are reported in Supplementary Table 1 of [36]. The variables included in the table include the patient's age, sex, clinical status (asymptomatic, upper respiratory infection or lower respiratory infection), duration of symptoms, and results of daily PCR testing for Days 0 to 6. Gautret et al. [36] report virological cure at Day 6 for 14 out of 20 patients treated with hydroxychloroquine and for 2 out of 16 in the control group, resulting in a p -value of 0.001 in an analysis not adjusted for any covariates.

The study by Gautret et al. [36] has been subject to criticism, mainly due to its limitations in design including the small sample size, choice of control patients, open label treatment and study discontinuations [5]. Although some preclinical data suggested potentially beneficial effects [26], there were also some early warnings regarding some potentially harmful effects [34]. In the meanwhile, data from large-scale randomized controlled trials are available demonstrating that hydroxychloroquine is not suitable for postexposure prophylaxis for or the treatment of COVID-19 [21,23,39]. The timeline of events is nicely depicted in Fig. 1 of a review by Sattui et al. [60].

3. Alternative analysis methods

3.1. The choice of effect measure

We consider a binary outcome Y as well as a binary treatment A (1: experimental treatment, 0: control) and a vector of observed covariates L . In clinical trial applications with a binary outcome Y as in our motivating example, logistic regression is usually the method of choice. This method of analysis experienced a huge increase in the 1980s [7] and is still very prevalent in clinical applications. The natural estimate obtained by a logistic regression is the odds ratio

$$OR = \frac{P(Y = 1|A = 1)/P(Y = 0|A = 1)}{P(Y = 1|A = 0)/P(Y = 0|A = 0)}, \quad (1)$$

i.e. the ratio of the odds of having the outcome under treatment and the odds of experiencing the outcome in the control group. The use of odds ratios, however, is not without criticism in the literature, see e. g. [19,25,27,63]. Common arguments against the use of the OR include that ORs are often not well understood by practitioners [30] or are misleadingly interpreted as relative risks [19], which is only appropriate with rare events. Other possible effect measures include the risk ratio or the risk difference. It is often argued that the risk difference is of greater importance to clinical decision makers than relative effect measures such as the OR [12]. Particularly in causal inference literature, there is a focus on the risk difference as effect measure. One reason for this is the issue of (non-)collapsibility: While marginal and conditional treatment effects coincide for the risk difference due to collapsibility, this is not true for the odds ratio [35,54,61]. The same arguments also hold for the hazard ratio obtained from a Cox model in case of time-to-event data. Thus, additive models are the preferred choice here as well [1,46].

3.2. Notation and some causal background

In a randomized controlled trial, one would assume that due to randomization, the influence of the covariates L is the same for treated and control patients. In observational studies, where allocation of the treatment is not in the hand of the investigator, this direct comparison of the treatments may no longer be fair due to the influence of other confounding factors, i.e., the distribution of the other risk factors L may differ between treated and controls. In order to imitate an RCT and to get valid estimates in this situation, a common approach is the so-called potential or counterfactual outcomes framework [38]: Let $Y^{a=1}$ denote the outcome that would have been observed under treatment value $a = 1$, and $Y^{a=0}$ the outcome that would have been observed under control ($a = 0$). A causal effect is now defined as follows: we say that A has a causal effect on Y if $Y^{a=1} \neq Y^{a=0}$ for an individual. In practice, however, only one of these outcomes is observed for an individual under study. Therefore, we can only ever estimate an *average causal effect*, which is present if $P(Y^{a=1} = 1) \neq P(Y^{a=0} = 1)$, i.e. the probability of the outcome under treatment is different from that under control, in the population of interest [38]. Thus, the causal risk difference is defined as

$$RD_c = P(Y^{a=1} = 1) - P(Y^{a=0} = 1) \quad (2)$$

3.3. Covariate adjustment of outcome model

The conventional method to correct for baseline differences between groups is adjusting for all relevant patient characteristics in the outcome regression model. To many medical statisticians, the natural choice of model for binary outcome data would be a logistic regression model. This, however, gives an estimate of the odds ratio, not the risk difference we are interested in. Moreover, in the case of small sample sizes, the number of possible variables to adjust for is limited by the observed events. Otherwise, logistic regression estimators may be biased or the model may not converge due to separation, i.e. a single covariate or a

combination of multiple covariates perfectly separates events from non-events [8,65,68,69]. Possibilities to correct for this include Firth's penalized logistic regression and extensions thereof, see [52] and the references cited therein. To obtain the risk difference, one could use a generalized linear model with Binomial distribution and identity link function [12]. However, the identity link function does not constrain the predicted probability to lie between 0 and 1 and the model often fails to converge [12,24]. An alternative, which avoids convergence issues, is to use ordinary least-squares estimation (OLS) instead, i.e. we assume a linear relationship

$$Y = \beta_0 + \beta_{tr}A + \alpha_1\ell_1 + \dots + \alpha_p\ell_p.$$

here, ℓ_1, \dots, ℓ_p denote the observed values of the covariates L_1, \dots, L_p . Although OLS is usually used to analyze the mean of a continuous outcome, it can also be used to estimate risk differences, since the mean is equal to the risk in case of a binary response coded as 0 and 1 [24]. Moreover, no distributional assumption is necessary to proof unbiasedness of the OLS estimator. In order to draw valid statistical inference, however, one has to consider robust variance estimators such as the Huber-White estimator. Since this is an asymptotic version of the robust variance, corrections are needed for small samples. The so-called HC3 variance estimator has been shown to perform best [24]. The idea is to multiply the Huber-White robust variance by a correction factor that converges to 1 as sample sizes increase. The HC3 variance estimator is available in R (package *sandwich*) and SAS (PROC *GLIMMIX*).

3.4. Propensity score based methods

Several different methods have been proposed to estimate RD_c in the literature, see e. g. [13,38] for an introduction. Many of these methods are based on the propensity score. The propensity score of individual i is defined as $\hat{p}_i := \hat{P}(A_i = 1|L_i)$, i.e., the estimated probability of receiving treatment given the covariates. For all methods considered in this paper, we estimate the propensity score using a logistic regression model for treatment allocation based on all observed covariates, i.e.

$$P(A = 1|L) = \frac{\exp(\beta_0 + \beta_1\ell_1 + \dots + \beta_p\ell_p)}{1 + \exp(\beta_0 + \beta_1\ell_1 + \dots + \beta_p\ell_p)}.$$

In a practical data analysis, there are several possibilities for taking the propensity score into account. We will describe the most common methods in the following and apply them to the data example.

PS covariate adjustment In this approach, the outcome Y is regressed on the estimated propensity score \hat{p} and the treatment exposure A , i.e., $Y = \beta_0 + \beta_{tr}A + \beta_{ps}\hat{p}$ and an estimator of the causal risk difference is given by $\hat{\beta}_{tr}$.

Matching on the propensity score Another possibility to balance treatment allocation is to match subjects on the propensity score. The idea is to find individuals with a similar propensity score in the treatment and the control group. There are various methods to match individuals. Particularly in small sample studies, it is impossible in practice to find exact matches. Thus, one needs to define an acceptable difference between the propensity scores of treated individuals and controls that will be used for matching. These differences are called *calipers* and should be small enough to allow for "a practical but meaningful equation of pairs" [6]. Following recently published recommendations [9], where propensity matching in small sample sizes was investigated, we performed a 1:1 nearest neighbor matching without replacement on the logit of the propensity score using calipers with a maximum width of 0.2 standard deviations. In this modification of classical nearest neighbor matching, subjects are only matched if the absolute difference of their propensity scores is within the pre-specified caliper distance [14]. This distance is usually defined as a proportion of the standard deviation of the logit of the propensity score. In R, this can e. g. be performed using the *MatchIt*-package, where the PS-model, the

method used for matching and the caliper can be specified. A caliper of 0.2 avoids matching dissimilar individuals. Note, however, that this setting differs from the default setting in R, where the caliper is set to 0.

In a matched cohort, we can calculate the risk difference as

$$(\bar{b} - \bar{c}) / n,$$

where \bar{b} is the number of pairs where the treated subject experiences the event whereas the untreated subject does not, \bar{c} are the pairs where the untreated subject experiences the event but the treated does not and n is the total number of matched pairs, see e. g. [12] for details.

Note that since we match individuals without replacement, the matched data set will usually be smaller than the original study, sometimes even discarding treated individuals.

Inverse probability of treatment weighting (IPTW) Inverse probability weighting uses the whole data set, but weighs each individual with his or her (inverse) probability of receiving the actually given treatment. This way, it generates a pseudo-population with (almost) perfect covariate balance between treatment groups. More specifically, IPTW assigns weight $w_i = 1/\hat{p}_i$ to treated subjects and weight $w_i = 1/(1 - \hat{p}_i)$ to controls. The resulting pseudo-population is analyzed using weighted regression with robust standard errors, which can, e. g. be obtained from the *survey*-package in R.

3.5. g-computation

The fourth possibility to account for covariate unbalance that we consider is *g-computation* [38,53], also known as *the parametric g-formula* or *direct standardization*, see [62] for an excellent introduction. The idea is that the marginal counterfactual risk can be written as

$$P(Y^a = 1) = \sum_{\ell} P(Y^a = 1|L = \ell)P(L = \ell) = \sum_{\ell} P(Y = 1|L = \ell, A = a)P(L = \ell)$$

here, the sum is over all values ℓ of the confounder(s) L that occur in the population. The right-hand side of this equation can now be estimated using the available data on Y , A and L . More precisely, we have to predict the potential outcome for every person i in the population assuming.

1. i was treated
2. i was a control

irrespective of the treatment actually received. In order to achieve this, we first fit a so-called Q-model to the data relating the outcome Y to the exposure A and to confounders L . For a binary outcome as in our situation, this is usually a logistic regression model. Instead of using this model for estimation of the treatment effect, however, we use it to predict $\hat{P}(Y = 1|L = \ell, A = 1)$ and $\hat{P}(Y = 1|L = \ell, A = 0)$ for all individuals by artificially creating two new data sets: One where $A = 1$ for all individuals and one where $A = 0$ for all individuals, respectively. Thus, this step can be thought of as imputing the missing potential outcomes for each subject in the population. Finally, the causal risk difference \widehat{RD}_c can be estimated by averaging over the estimated probabilities of the outcome under treatment and control and applying Eq. (2).

Confidence intervals for *g-computation* are usually obtained by a nonparametric bootstrap approach [29,38], i.e. by drawing with replacement from the data and analyzing each bootstrap data set like we analyzed the original data. Resampling approaches like this lead to asymptotically valid inference procedures [29] and have been shown to be superior in small samples in various situations [18,20,32,33,42,50]. The number of bootstrap repetitions should be chosen reasonably large.

Table 1

Association to treatment assignment and outcome used in the simulation scenario motivated by Austin [11] (Scenario 3).

	Strongly associated with treatment	Moderately associated with treatment	Not associated with treatment
Strongly associated with outcome	x_1	x_2	x_3
Moderately associated with outcome	x_4	x_5	x_6
Not associated with outcome	x_7	x_8	x_9

Table 2

Overview of the simulated scenarios and where to find the results.

	β_{trt}	true RD	β_0	Percent treated on average	simulated for true RD
Scenario 1	0	0	-2.3	55%	0, 0.16, 0.4
	0.8678	0.16	-1.8	66%	0
	3.128	0.4	-1	80%	0
Scenario 2	0	0	-2.3	54%	0, 0.16, 0.4
	1.1111	0.16	-2.3	54%	0, 0.16, 0.4
	3.71	0.4	-2.3	54%	0, 0.16, 0.4
Scenario 3	0	0	-3.5	49%	0, 0.16, 0.4
	1.032	0.16	-1.5	80%	0
	2.448	0.4	-1.5	80%	0
Results:	Figs. 1 and 2, Tables 3-5		Figs. 3 and 4, Table 6		

We used 1000 bootstrap repetitions in the simulation study, but recommend a higher number in real-life applications. Upper and lower 95% confidence intervals are obtained using the 2.5 and 97.5 percentiles of the bootstrap distribution. Note that a statistical test can be obtained similarly by calculating the test statistic in each bootstrap sample and then comparing the original test statistic to the empirical $(1 - \alpha)$ -quantile of the bootstrap distribution. A p -value is obtained by counting how often the original test statistic is smaller than the bootstrap statistic and dividing this number by the conducted bootstrap replications, see e. g. [32,33] for similar approaches. To investigate the small sample performance of the bootstrap for *g-computation* in detail and determine whether more elaborate bootstrap techniques might lead to better performance is part of future research.

3.6. Doubly robust estimators

While IP weighting requires the propensity model to be correct, i.e. a correct model for the treatment A conditional on confounders L , the *g-formula* requires a correct model for the outcome Y conditional on treatment A and the confounders L , the Q-model. A doubly robust (DR) estimator, in contrast, is consistent if at least one of the two models is correctly specified. There are many types of doubly robust estimators (see e. g. [16,41,67] and the references cited therein for an overview). We will focus on three different ones here. The first two are applied to the *g-computation* whereas the third is an extension of IPTW.

Simple DR g-computation The first DR estimator we consider is a very simple one [16,38]: First, we estimate the weights w_i as described above. We then fit our Q-model to the data including an additional covariate z , where $z_i = w_i$ if $A_i = 1$ and $z_i = -w_i$ if $A_i = 0$. Finally, we again obtain a causal risk difference from Eq. (2). This method is referred to as “Simple DR *g-computation*” in the following. Kang and Schafer [41] studied the performance of different DR estimators with a particular focus on the situation, where both the outcome and the PS model are misspecified. They found that this estimator behaves poorly, when the PS-model is misspecified and even state that “[t]he performance of this method is disastrous when some of the estimated propensities are small”

[41].

DR using quintiles Another possibility for a DR estimator also studied by Kang and Schafer and found to “[perform] better than any of the other DR methods when the [models] are both incorrect” [41] is obtained by coarsening the logit of the estimated propensity score into five categories according to the quintiles. Thus, we include four dummy variables distinguishing among these categories in the Q-model for the g-computation, see [41] for details. We denote this approach “DR using quintiles” in the following.

Augmented IPW (AIPW) Another approach is to augment the IPTW estimator described above with a regression model for the outcome variable. Thus, a separate outcome model of Y on the confounders L is needed. Details on the method can be found in [37,45] and the resulting AIPW estimator is implemented in the R package PSW. Note that this approach is closely connected to the simple DR g-computation described above: Bang and Robins [16] found that the augmented estimator can be viewed as an unweighted regression including the inverse of the PS as a covariate [37].

4. Simulation study

The set-up of our simulation study closely followed Austin [11,12]. The data-generating process is as follows: First, we generate n covariates x_1, \dots, x_n (see the following subsections for details). We then generate the treatment status for each subject $i = 1, \dots, N$ according to the model

$$\text{logit}(p_{i,\text{treatment}}) = \beta_0 + \beta_1 x_1 + \dots + \beta_n x_n. \tag{3}$$

Treatment is then randomly assigned to each subject following a

Bernoulli distribution with subject-specific probability of treatment assignment $A_i \sim \text{Bernoulli}(p_{i,\text{treatment}})$. Next, the outcome Y_i of each subject is simulated conditional on treatment assignment A_i and the covariates associated with the outcome according to

$$\text{logit}(p_{i,\text{outcome}}) = \alpha_0 + \beta_{\text{trt}} A_i + \alpha_1 x_1 + \dots + \alpha_n x_n \tag{4}$$

and $Y_i \sim \text{Bernoulli}(p_{i,\text{outcome}})$. Here, β_{trt} denotes the log-odds ratio relating treatment to the outcome. Thus, a value of $\beta_{\text{trt}} = 0$ corresponds to the null effect, i.e. an odds ratio of 1 and a risk difference of 0.

In contrast to odds ratios, the risk difference is collapsible, i.e. the average subject-specific risk difference is equal to the marginal risk difference. Based on 1000 data sets of size $N = 10,000$, we used the following procedure to determine the average risk difference and adjust the value of β_{trt} to obtain the desired non-null risk differences: For a fixed value of β_{trt} , we generate the counterfactual outcomes under treatment ($A_i = 1$) and control ($A_i = 0$) for each individual and calculate the marginal probabilities under treatment and control. The risk difference is then equal to the difference between these two marginal probabilities [12]. Using an iterative process, we modified β_{trt} until we got close enough to the desired marginal risk difference.

Concerning the covariates, we considered three different scenarios:

4.1. Scenario 1: the COVID-19 example

The first scenario aimed at mimicking the data example. Thus, we generated four covariates:

Table 3

Median length of the 95% confidence intervals (Length CI), root mean square error of the estimated treatment effect (RMSE), median absolute error (MAE) and number of models failed for Scenario 1: COVID-19. Crude RD is estimated by a linear regression adjusting only for treatment assignment, Cov adjustment additionally adjusts for baseline covariates, PS covariate denotes the method including the PS in the outcome regression model and Simple DR g-comp and DR using quintiles refer to the doubly robust g-computation methods, respectively.

true RD	N		Crude RD	covariate adjusted	PS cov	matching	IPTW	Simple DR g-comp	DR using quintiles	AIPW
0	N = 40	CI Length	0.6442	0.7157	0.736	4.162	0.2569	1.15	0.8127	0.5355
		RMSE	0.2105	0.169	0.1722	0.2185	0.1847	0.39	0.1886	0.2662
		MAE	0.15	0.1133	0.1142	0.1538	0.1246	0.3695	0.1265	0.2656
		Failures	0	0	0	1404	0	1	0	0
	N = 100	CI Length	0.397	0.4072	0.4251	4.219	0.1463	1.001	0.4006	0.5371
		RMSE	0.1621	0.09911	0.09948	0.1206	0.103	0.3388	0.1001	0.2497
		MAE	0.1319	0.06541	0.06507	0.08696	0.06772	0.2904	0.06678	0.2493
		Failures	0	0	0	1647	0	0	0	0
	N = 1000	CI Length	0.1236	0.121	0.1265	4.099	0.04152	0.5677	0.1179	0.5354
		RMSE	0.1319	0.0304	0.03046	0.06247	0.03058	0.1609	0.03002	0.2636
		MAE	0.1278	0.02041	0.02068	0.0596	0.02063	0.1054	0.02028	0.2632
		Failures	0	0	0	1998	0	0	0	0
0.16	N = 40	CI Length	0.6107	0.6981	0.7196	4.655	0.247	1.15	0.8185	0.5656
		RMSE	0.1986	0.1638	0.1665	0.2076	0.1783	0.3857	0.1851	0.1819
		MAE	0.14	0.1113	0.113	0.1538	0.1219	0.375	0.1271	0.1228
		Failures	0	0	0	1404	0	0	0	58
	N = 100	CI Length	0.3772	0.3988	0.4164	4.773	0.1431	0.9995	0.3957	0.3652
		RMSE	0.1534	0.09643	0.09675	0.1154	0.1013	0.3292	0.09841	0.0999
		MAE	0.1265	0.06524	0.06522	0.07529	0.0679	0.3021	0.06641	0.06872
		Failures	0	0	0	1647	0	0	0	0
	N = 1000	CI Length	0.1174	0.1187	0.1245	4.651	0.04004	0.5618	0.1162	0.1181
		RMSE	0.1245	0.03063	0.03058	0.07065	0.03059	0.1592	0.02999	0.03044
		MAE	0.1213	0.02041	0.02066	0.07063	0.0201	0.1128	0.01986	0.02003
		Failures	0	0	0	1998	0	0	0	0
0.4	N = 40	CI Length	0.5111	0.6016	0.6038	5.227	0.1897	1.15	0.75	0.487
		RMSE	0.1509	0.1453	0.148	0.1679	0.159	0.379	0.1671	0.1594
		MAE	0.1092	0.09989	0.1012	0.1	0.1082	0.325	0.1092	0.1007
		Failures	0	0	0	1366	0	0	0	61
	N = 100	CI Length	0.3169	0.3422	0.3493	5.115	0.1158	0.9894	0.351	0.3153
		RMSE	0.1177	0.08743	0.08763	0.09178	0.0911	0.3206	0.09038	0.09018
		MAE	0.09127	0.05963	0.05935	0.06429	0.06132	0.266	0.05985	0.06143
		Failures	0	0	0	1675	0	0	0	0
	N = 1000	CI Length	0.09881	0.1024	0.1054	4.978	0.03402	0.4898	0.1018	0.1028
		RMSE	0.08764	0.02848	0.02804	0.03864	0.02594	0.1442	0.02565	0.02599
		MAE	0.08488	0.02039	0.01987	0.04186	0.01707	0.09828	0.01714	0.0171
		Failures	0	0	0	1997	0	0	0	0

1. x_1 (representing sex) followed a Bernoulli distribution with parameter 0.5
2. x_2 (representing age) was drawn from a $N(45, 15)$ distribution and rounded to integers
3. x_3 (clinical status) was simulated as a categorical covariate with three categories, i.e. a $\text{Bin}(2, 0.5)$ distribution
4. x_4 (time since onset of disease) was generated from a uniform distribution on $[0, 10]$ and rounded to integers.

Treatment status was then generated according to Eq. (3) with

$$(\beta_0, \beta_1, \beta_2, \beta_{3,1}, \beta_{3,2}, \beta_4) = (-2.3, 0.31, 0.03, 1.099, -0.1054, 0.1031).$$

here, $\beta_{3,1}$ and $\beta_{3,2}$ correspond to the dummy-coded categories $x_{3,1}$ and $x_{3,2}$ for $x_3 = 1$ and $x_3 = 2$, respectively. The parameters were obtained from the data by univariate logistic regression. Note that this implies a moderate association of treatment with x_1 , $x_{3,2}$ and x_4 , a weak association with x_2 and a strong association with $x_{3,1}$.

Similarly, the outcome was generated following Eq. (4) with

$$(\alpha_0, \alpha_1, \alpha_2, \alpha_{3,1}, \alpha_{3,2}, \alpha_4) = (-1.06, 0.619, 0.0077, 0.9461, -1.3499, 0.0896),$$

implying a moderate association with x_1 and x_4 , a strong association with x_3 and a weak association with x_2 . The parameter β_{trt} was varied to generate different risk differences in the following way: For $\beta_{trt} = 0$, the risk difference is equal to 0. For $\beta_{trt} = 0.8678$ we get a risk difference of 0.16 and for $\beta_{trt} = 3.128$ the true risk difference equals 0.4. Finally, $\beta_0 = -2.3$ resulted in a similar distribution of treated individuals and controls as in the original data, yielding an average of 55.21% of individuals

in the treatment group. To study the influence of more or less unbalanced treatment groups, we also varied this parameter in the simulations. In particular, we additionally considered a treatment allocation of approx. 2:1 and 4:1.

4.2. Scenario 2: unmeasured confounder

The parameters in this setting are identical to Scenario 1, but we additionally added an unmeasured confounder. Thus, we simulated a covariate x_5 following an $N(0, 1)$ distribution with a strong effect on both treatment assignment and outcome. Therefore, β_5 and α_5 were set to log(5). However, x_5 entered neither the propensity score model nor the Q-model for the g-computation. For a risk difference of 0.16 and 0.4, β_{trt} was set to 1.1111 and 3.71, respectively.

4.3. Scenario 3: following Austin's design

This scenario is based on Austin [11]. Therefore, we used the same set-up as he did, namely simulating 9 binary covariates with different association to treatment assignment and outcome as described in Table 1.

Here, a strong association is represented by a coefficient of $\log(5)$, i.e. $\beta_1 = \beta_4 = \beta_7 = \alpha_1 = \alpha_2 = \alpha_3 = \log(5)$, while a moderate association has a coefficient of $\log(2)$, i.e. $\beta_2 = \beta_5 = \beta_8 = \alpha_4 = \alpha_5 = \alpha_6 = \log(2)$. We chose $\beta_0 = -3.5$ to obtain a balanced design with respect to treatment and α_0 was set to -5 . For more details on the simulation set-up, see [11]. The propensity score model and the Q-model included all 9 covariates. For a risk difference of 0.16 and 0.4, β_{trt} was set to 1.032 and 2.448,

Table 4

Median length of the 95% confidence intervals (Length CI), root mean square error of the estimated treatment effect (RMSE), median absolute error (MAE) and number of models failed for Scenario 2: Unmeasured confounder. Crude RD is estimated by a linear regression adjusting only for treatment assignment, Cov adjustment additionally adjusts for baseline covariates, PS covariate denotes the method including the PS in the outcome regression model and Simple DR g-comp and DR using quintiles refer to the doubly robust g-computation methods, respectively.

true RD	N		Crude RD	covariate adjusted	PS cov	matching	IPTW	Simple DR g-comp	DR using quintiles	AIPW
0	N = 40	CI Length	0.6118	0.7142	0.7028	3.566	0.2116	1.15	0.7985	0.5803
		RMSE	0.3781	0.3417	0.3416	0.363	0.3483	0.4579	0.3475	0.347
		MAE	0.3485	0.297	0.3005	0.3077	0.3031	0.3741	0.3032	0.2975
		Failures	0	0	0	1255	0	0	0	47
	N = 100	CI Length	0.3761	0.4038	0.4083	3.732	0.117	1.056	0.4009	0.3738
		RMSE	0.3581	0.3143	0.3144	0.3193	0.3147	0.4389	0.3152	0.3146
		MAE	0.3474	0.3032	0.3033	0.3043	0.3004	0.3504	0.3002	0.3005
		Failures	0	0	0	1502	0	0	0	0
	N = 1000	CI Length	0.1169	0.1204	0.1233	3.768	0.03334	0.7931	0.1197	0.1204
		RMSE	0.3471	0.2988	0.2989	0.3083	0.2987	0.3483	0.2983	0.2986
		MAE	0.346	0.2981	0.298	0.301	0.2978	0.2782	0.2976	0.2979
		Failures	0	0	0	1985	0	0	0	0
0.16	N = 40	CI Length	0.5639	0.6632	0.6563	3.823	0.1863	1.15	0.7525	0.5399
		RMSE	0.3553	0.3243	0.3247	0.3372	0.3314	0.4361	0.3253	0.3297
		MAE	0.329	0.2889	0.2901	0.2945	0.2891	0.41	0.2851	0.2903
		Failures	0	0	0	1255	0	0	0	47
	N = 100	CI Length	0.3454	0.3761	0.3791	3.837	0.103	1.059	0.3729	0.3485
		RMSE	0.3419	0.3056	0.3056	0.3111	0.3046	0.4174	0.3057	0.3046
		MAE	0.3304	0.2902	0.2904	0.2958	0.2903	0.3981	0.2901	0.2913
		Failures	0	0	0	1502	0	0	0	0
	N = 1000	CI Length	0.1075	0.1125	0.1149	3.85	0.02938	0.801	0.1123	0.1127
		RMSE	0.3295	0.2888	0.2888	0.2885	0.2875	0.3276	0.2871	0.2874
		MAE	0.3284	0.2881	0.288	0.2889	0.2863	0.2847	0.2866	0.2861
		Failures	0	0	0	1985	0	0	0	0
0.4	N = 40	CI Length	0.4687	0.5579	0.546	3.859	0.1567	1.15	0.5525	0.4621
		RMSE	0.2642	0.2481	0.2486	0.2593	0.2533	0.4077	0.2475	0.2559
		MAE	0.2429	0.2132	0.211	0.2154	0.2145	0.4119	0.2134	0.2149
		Failures	0	1	0	1278	0	0	0	51
	N = 100	CI Length	0.2908	0.3193	0.3185	3.815	0.0913	1.039	0.3235	0.2982
		RMSE	0.242	0.2187	0.2186	0.2212	0.218	0.363	0.2167	0.2178
		MAE	0.2303	0.2034	0.2036	0.2111	0.2044	0.3615	0.2009	0.203
		Failures	0	0	0	1513	0	0	0	0
	N = 1000	CI Length	0.09057	0.09519	0.09628	3.632	0.02705	0.7431	0.09586	0.09596
		RMSE	0.2335	0.2077	0.2076	0.2197	0.2046	0.2573	0.2042	0.2046
		MAE	0.2327	0.207	0.2068	0.2274	0.2032	0.2168	0.2032	0.2031
		Failures	0	0	0	1990	0	0	0	0

Table 5

Median length of the 95% confidence intervals (Length CI), root mean square error of the estimated treatment effect (RMSE), median absolute error (MAE) and number of models failed for Scenario 3: Austin. Crude RD is estimated by a linear regression adjusting only for treatment assignment, Cov adjustment additionally adjusts for baseline covariates, PS covariate denotes the method including the PS in the outcome regression model and Simple DR g-comp and DR using quintiles refer to the doubly robust g-computation methods, respectively.

true RD	N		Crude RD	covariate adjusted	PS cov	matching	IPTW	Simple DR g-comp	DR using quintiles	AIPW
0	N = 40	CI Length	0.5603	0.762	0.7935	1.715	0.1818	1.125	0.925	1.126
		RMSE	0.194	0.1726	0.1945	0.2163	0.1973	0.351	0.235	0.2144
		MAE	0.15	0.1132	0.1211	0.1429	0.1343	0.3	0.2	0.1389
		Failures	0	0	0	982	0	0	0	272
	N = 100	CI Length	0.3439	0.3961	0.4331	1.887	0.1601	0.8337	0.4516	0.4447
		RMSE	0.1592	0.09837	0.1008	0.1221	0.1233	0.2284	0.1036	0.119
		MAE	0.1345	0.06787	0.06938	0.08696	0.08036	0.1498	0.07362	0.07635
		Failures	0	0	0	788	0	0	0	0
	N = 1000	CI Length	0.1068	0.114	0.1267	1.939	0.06029	0.2382	0.1097	0.1366
		RMSE	0.1356	0.02942	0.02966	0.04048	0.03514	0.06127	0.02884	0.03374
		MAE	0.133	0.01964	0.01961	0.02799	0.02342	0.04113	0.01954	0.02237
		Failures	0	0	0	600	0	0	0	0
0.16	N = 40	CI Length	0.5831	0.8028	0.8645	2.195	0.2533	1.125	0.9497	1.074
		RMSE	0.2098	0.1824	0.2139	0.2327	0.2145	0.3532	0.2429	0.2261
		MAE	0.1582	0.1217	0.131	0.16	0.1505	0.3209	0.1789	0.1484
		Failures	0	0	0	982	0	0	0	272
	N = 100	CI Length	0.3582	0.4209	0.471	2.42	0.2227	0.8592	0.4589	0.4449
		RMSE	0.1764	0.1033	0.1059	0.1281	0.138	0.2297	0.1103	0.1286
		MAE	0.1533	0.06929	0.07238	0.08593	0.08858	0.1635	0.07497	0.08366
		Failures	0	0	0	788	0	0	0	0
	N = 1000	CI Length	0.1115	0.1214	0.1378	2.472	0.07675	0.2576	0.1191	0.1444
		RMSE	0.1529	0.03138	0.03129	0.0446	0.03796	0.06673	0.03052	0.03651
		MAE	0.1505	0.02178	0.02165	0.0316	0.02533	0.043	0.02084	0.02467
		Failures	0	0	0	584	0	0	0	0
0.4	N = 40	CI Length	0.5439	0.7943	0.867	2.695	0.2462	1.125	0.925	0.9721
		RMSE	0.1965	0.1781	0.2048	0.2196	0.2123	0.3738	0.2159	0.2215
		MAE	0.1514	0.1198	0.1323	0.15	0.1578	0.3061	0.1277	0.1423
		Failures	0	0	0	982	0	0	0	272
	N = 100	CI Length	0.3342	0.4202	0.4732	2.846	0.2415	0.9199	0.4436	0.4277
		RMSE	0.1692	0.1036	0.106	0.1225	0.1442	0.2462	0.1115	0.1276
		MAE	0.1484	0.06843	0.0715	0.08	0.09595	0.1785	0.07564	0.08568
		Failures	0	0	0	788	0	0	0	0
	N = 1000	CI Length	0.1036	0.1219	0.1387	2.863	0.09316	0.2694	0.1262	0.1481
		RMSE	0.1486	0.03536	0.03439	0.04698	0.04086	0.06804	0.03299	0.03778
		MAE	0.1463	0.02417	0.02408	0.03327	0.02782	0.04509	0.02271	0.02512
		Failures	0	0	0	584	0	0	0	0

respectively. In addition to Austin’s setting with an equal treatment allocation of 1:1, we also considered a situation with approx. 4:1 treated patients.

An overview of all simulated scenarios is given in Table 2.

In order to compare our results for the risk difference to the approach of a logistic regression, i.e. to estimating a causal odds ratio, we have also performed our simulations for the odds ratio. The results are included in the Supplemental material.

4.4. Simulation results

To study the influence of small sample sizes on the methods, we simulated N = 40,100,1000 individuals for each scenario. Simulations were performed in R Version 3.6.3 with 2000 simulation runs and the bootstrap confidence intervals for the g-computation are based on 1000 bootstrap replications. Note that while 1000 bootstrap replications suffice in simulations, we would recommend a higher number, say 10,000, in real-life applications.

We used different measures to compare the results. With respect to the point estimators, we considered the mean bias, i.e. the mean difference between the true risk difference RD and the estimated risk difference \widehat{RD} . The results are displayed in Fig. 1. Moreover, the root mean square error of each estimated risk difference (RMSE) and the median of the absolute errors (MAE), i.e. the median of $|\widehat{RD} - RD|$ are displayed in Tables 3–5.

Concerning the confidence intervals, we considered the percentage of 95% confidence intervals that contained the true risk difference (coverage probability) as well as the median length of the 95%

confidence interval. These measures are displayed in Fig. 2 and Tables 3–5, respectively. Finally, we also reported how often the methods failed, e. g. since no matching could be performed or the model did not converge. These were excluded from the calculations and reported as failures in the tables.

For comparison, we included the crude as well as the covariate-adjusted risk difference.

Across all scenarios considered here, we note that the matching procedure is the most prone to failure. Even for the large sample sizes, it often fails in creating a matched sample. This is even more pronounced for the situations with unbalanced treatment allocation, see Table 6. These results are in line with the findings of Adrillon et al. [9], who stress the need for development of appropriate matching methods in small sample studies. Furthermore, we found that using the default caliper, which is 0 in R, leads to extremely biased results with coverage probabilities dropping below 1% in some situations (results not shown).

We note that the mean bias of all methods decreases with growing sample sizes, although the difference is not pronounced. The largest mean bias is observed for the crude RD estimation. For Scenario 1 with a true risk difference of 0, however, the AIPW method has the largest bias, see Fig. 1. Our simulations also show very good results for simple covariate adjustment with respect to both RMSE and MAE.

With respect to coverage, we observe surprisingly poor coverage probabilities for IPTW (Fig. 2) and at the same time very short confidence intervals (Tables 3–5). For the small sample sizes, the coverage of IPTW is even worse than the crude risk difference. The other methods show similar results except for AIPW, which again cannot handle Scenario 1 for a risk difference of 0 very well.

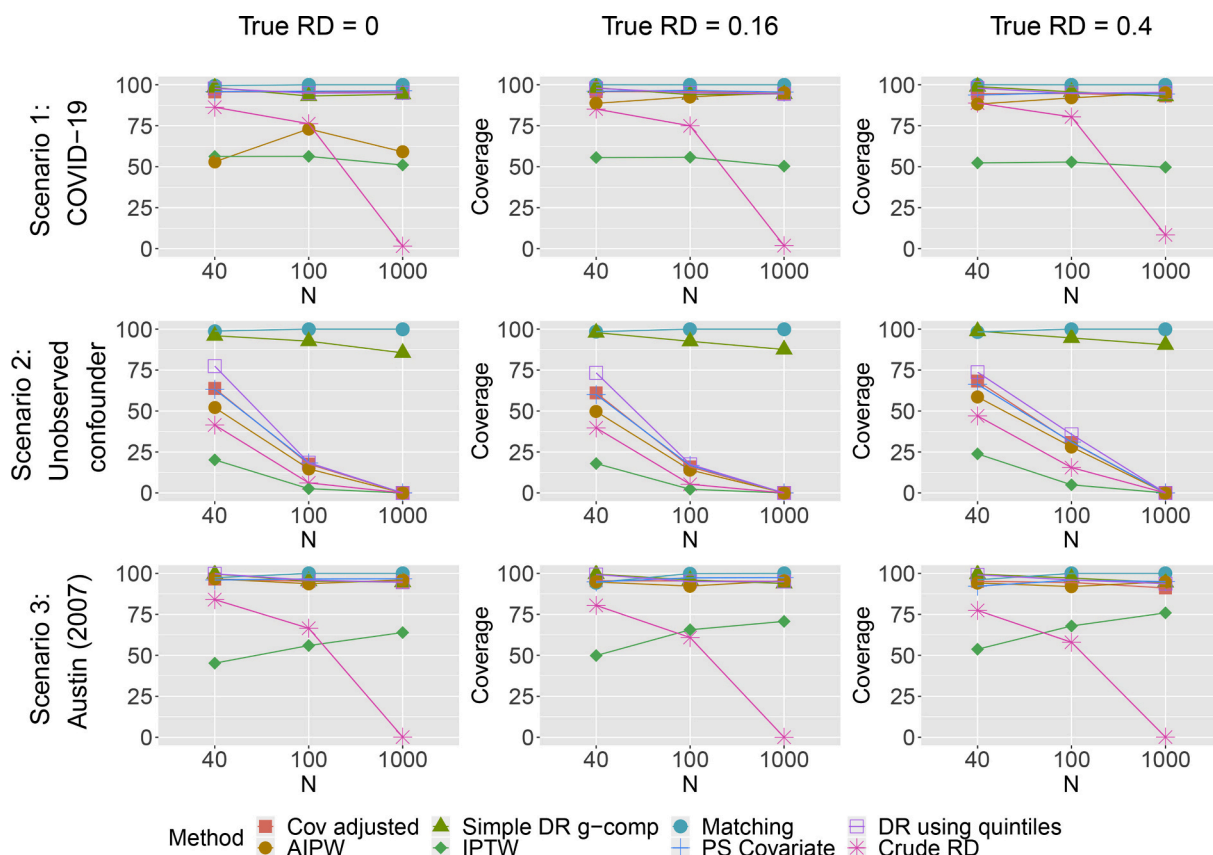


Fig. 2. Displayed is the coverage probability (in %) for the three scenarios (rows) and the three simulated risk differences (columns).

As expected, the results observed for Scenario 2 show a larger bias and much smaller coverage probabilities than in the other scenarios due to the unobserved confounder. Here, the simple DR g-computation performs best both with respect to bias and coverage probabilities. Coverage probabilities are also very high for matching, but due to the many failures and the extremely wide confidence intervals of the method, these results should be interpreted with caution.

Figs. 3 and 4 show the mean bias and coverage probabilities of the methods in Scenario 1 and 3, respectively, where we varied the proportion of individuals who receive treatment. The results are very similar to the ones observed for the balanced situation.

Concerning the doubly robust estimators, we find that the DR using quintiles performs best across Scenarios 1 and 3. Only in Scenario 2, the simple DR g-computation shows the best results. However, more research is needed here to investigate whether this stems from different biases in opposite directions which compensate each other.

A few comments on the comparison between OR and RD are in place: As can be seen from the results in the supplemental material, the logistic regression is very unstable for small sample sizes. Thus, we observe a lot more failures than we did for the RD, where only matching and (for small samples) AIPW showed relevant failures. Moreover, the estimation of the OR is sometimes heavily biased, resulting in a large mean bias as opposed to a relatively mild median bias, which in turn also leads to huge confidence intervals.

4.5. Recommendations for small-scale studies

Based on the simulation results, we deduct the following recommendations for applications in clinical studies:

1. Causal inference methods can correct for the *non-randomized* nature of a study. However, they cannot deal with other issues such as flaws

in the study design, data quality etc. This is nicely demonstrated by the simulation results observed for Scenario 2: The presence of an unmeasured confounder renders the methods practically useless. Thus, to speak with Rubin’s words, the most important recommendation is: “For objective causal inference, design trumps analysis” [58].

2. In small sample settings, the risk difference provides a more stable effect measure than the odds ratio, which is due to the limitations of the logistic regression in small samples. Thus, the risk difference is the preferred effect measure in small samples. This recommendation does not only apply to the causal inference methods but also to covariate adjustment.
3. For small total sample sizes ($N = 40$), the best performance was observed for covariate adjustment, PS covariate and the DR using quintiles, i.e. a doubly robust g-computation.
4. IPTW performed well with respect to bias, RMSE and MAE, but due to its extremely low coverage probability it cannot be recommended.
5. Since one can always only investigate a limited number of simulated settings, we recommend conducting simulations for a given example at hand. In order to facilitate this, the R code used for the simulations in this paper is available from Github (https://github.com/smn74/CIM_COVID-19).

5. Discussion

In ongoing pandemics there is an urgent unmet medical need to develop vaccines, diagnostics and treatments in a very timely fashion. Despite the time pressure, however, the standards of evidence should not unduly be lowered [17,31,49,55]. Using a small-scale non-randomized study in COVID-19 [36] as a motivating example, we discuss how robust analyses can be conducted by use of appropriate causal inference methods.

Table 6

Median length of the 95% CIs (Length CI), RMSE, MAE and number of failures for Scenario 1: COVID-19 and Scenario 3: Austin with unbalanced treatment allocation.

Scenario	N		Crude RD	covariate adjusted	PS cov	matching	IPTW	Simple DR g-comp	DR using quintiles	AIPW
Scenario 1 66% treated	N = 40	CI Length	0.6684	0.752	0.7806	4.074	0.3096	1.15	0.8421	0.6278
		RMSE	0.2136	0.176	0.1795	0.2295	0.1987	0.3823	0.1961	0.2544
		MAE	0.1535	0.117	0.1182	0.125	0.1331	0.3484	0.1321	0.1307
		failures	0	1	0	1936	0	1	0	60
	N = 100	CI Length	0.4158	0.4264	0.4463	.	0.1949	0.9479	0.4172	0.3925
		RMSE	0.1674	0.1054	0.1061	.	0.114	0.3129	0.1067	0.1134
		MAE	0.1298	0.06967	0.0703	.	0.07415	0.2511	0.07192	0.07421
		failures	0	0	0	2000	0	0	0	0
	N = 1000	CI Length	0.1298	0.1258	0.1323	.	0.05847	0.4306	0.1228	0.1277
		RMSE	0.1333	0.0322	0.03225	.	0.03336	0.1177	0.03158	0.03311
		MAE	0.129	0.02163	0.02188	.	0.02199	0.07594	0.0214	0.02202
		failures	0	0	0	2000	0	0	0	0
Scenario 1 80% treated	N = 40	CI Length	0.8139	0.8853	0.9398	.	0.4916	1.171	0.9248	0.8298
		RMSE	0.2454	0.2088	0.2178	.	0.2577	0.3901	0.2345	0.8889
		MAE	0.1714	0.1376	0.1448	.	0.1853	0.3552	0.16	0.1929
		failures	1	2	1	2000	1	0	0	605
	N = 100	CI Length	0.4956	0.4942	0.526	.	0.4647	0.9358	0.4874	0.4724
		RMSE	0.1761	0.1218	0.1239	.	0.1516	0.2775	0.1245	0.1569
		MAE	0.1294	0.08333	0.08371	.	0.1018	0.2026	0.08554	0.09977
		failures	0	0	0	2000	0	0	0	7
	N = 1000	CI Length	0.1531	0.1458	0.1546	.	0.1476	0.3277	0.143	0.1599
		RMSE	0.1365	0.03668	0.03673	.	0.0409	0.08494	0.03639	0.04032
		MAE	0.1319	0.02492	0.02469	.	0.02813	0.05294	0.02466	0.02801
		failures	0	0	0	2000	0	0	0	0
Scenario 3 80% treated	N = 40	CI Length	0.6086	0.8389	0.8496	.	0.1076	1.15	0.975	1.829
		RMSE	0.1992	0.1929	0.2359	.	0.2228	0.4026	0.2751	0.4225
		MAE	0.1562	0.1223	0.1352	.	0.1785	0.35	0.225	0.2219
		failures	3	3	3	2000	3	0	0	1690
	N = 100	CI Length	0.3744	0.4348	0.4804	.	0.2419	0.8949	0.5733	0.9399
		RMSE	0.1636	0.1029	0.1092	.	0.2022	0.2403	0.1412	0.2176
		MAE	0.1382	0.06706	0.07216	.	0.1492	0.1709	0.09435	0.1439
		failures	0	0	0	2000	0	0	0	45
	N = 1000	CI Length	0.1166	0.1226	0.136	.	0.439	0.2122	0.146	0.2318
		RMSE	0.1338	0.03366	0.03193	.	0.07436	0.05153	0.03761	0.06406
		MAE	0.1304	0.02278	0.0222	.	0.04846	0.03508	0.02573	0.04318
		failures	0	0	0	2000	0	0	0	0

The conventional method to correct for baseline differences between groups is adjusting for all relevant patient characteristics in the outcome regression model. This is, however, not favorable for different reasons. As Rosenbaum and Rubin [56] point out, covariate adjustment works poorly in cases where e. g. the variance of a covariate is unequal in the treatment and the control group. A commonly applied alternative in observational studies are propensity score methods. Since these methods were derived from a formal model for causal inference, their use allows for well-defined causal questions [37,47]. Moreover, propensity score methods also work as a dimension reduction tool by combining multiple covariates into a single score [47,57]. This is especially important in situations with a large number of covariates compared to the number of subjects under study. Different approaches have been suggested for the PS modeling strategy. Originally, nonparsimonious models including all potential confounders have been recommended for the propensity score [59]. This approach, however, may not be feasible in small samples. Thus, it has been recommended [9,22,51] to use some kind of variable selection procedure in this case, but clear recommendations are lacking [37]. Importantly, the choice of the variables should not be based on some goodness-of-fit measure [22,72] but rather on the relationship of the variables with both the outcome and the exposure [51]. Since our data example only contained four potential confounders, we have included them all in the PS model and have not investigated methods of variable selection here.

In line with [41] our simulation studies showed that different DR estimators led to different results in the scenarios considered. More thorough investigations on this topic, especially concerning the type of DR adjustment, will be part of future research.

Some comments on the estimands obtained by the different methods are in place: First, our aim was to estimate the average causal effect in our study population. When there is a lack of overlap between the propensity

score distributions of the two groups, a problem quite common for small samples, IPTW may become unstable due to extremely large weights. Mao et al. [45] recently proposed modified weights, which result in a different estimand that deviates from the average treatment effect. However, this approach is appropriate in treatment effect discovery, which is often the main motivation of small observational studies. Similarly, propensity score matching creates a population where treated individuals, who cannot be matched to any control patients, are excluded. Thus, the effect estimate obtained here corresponds to a subset of the population, which is hard to describe. Since the matched population is not very well characterized, it is difficult to generalize results obtained there to the general population [38]. Moreover, PS matching is also criticized for the fact that a large number of irrelevant covariates might lead to matched pairs which actually differ in relevant covariates [28,71]. Solutions to this problem use machine learning techniques to first determine the relevant covariates and then match exactly on these [28,71]. Since they require large training and matching sets, however, they cannot be applied to small samples. Furthermore, we did not account for the fact that the propensity score used for matching is itself estimated, see [4] for a thorough discussion of this topic. Moreover, it should be noted that classical bootstrap approaches such as the nonparametric bootstrap we used in the g-computation are not applicable to matching estimators [2,3]. It is also worth noting that among the methods we discussed here, only IPT weighting and g-computation can be generalized to more complex situations involving time-varying treatments [38]. Finally, it has to be noted that when estimating the odds ratio instead of the risk difference, *marginal* and *conditional* treatment effects differ due to non-collapsibility of the odds ratio [15,44]. It should be noted that covariate adjusted logistic regression, PS covariate adjusted logistic regression, and conditional logistic regression in the matched sample all estimate the conditional OR and not the marginal OR in this case. Thus, care has to be taken when comparing

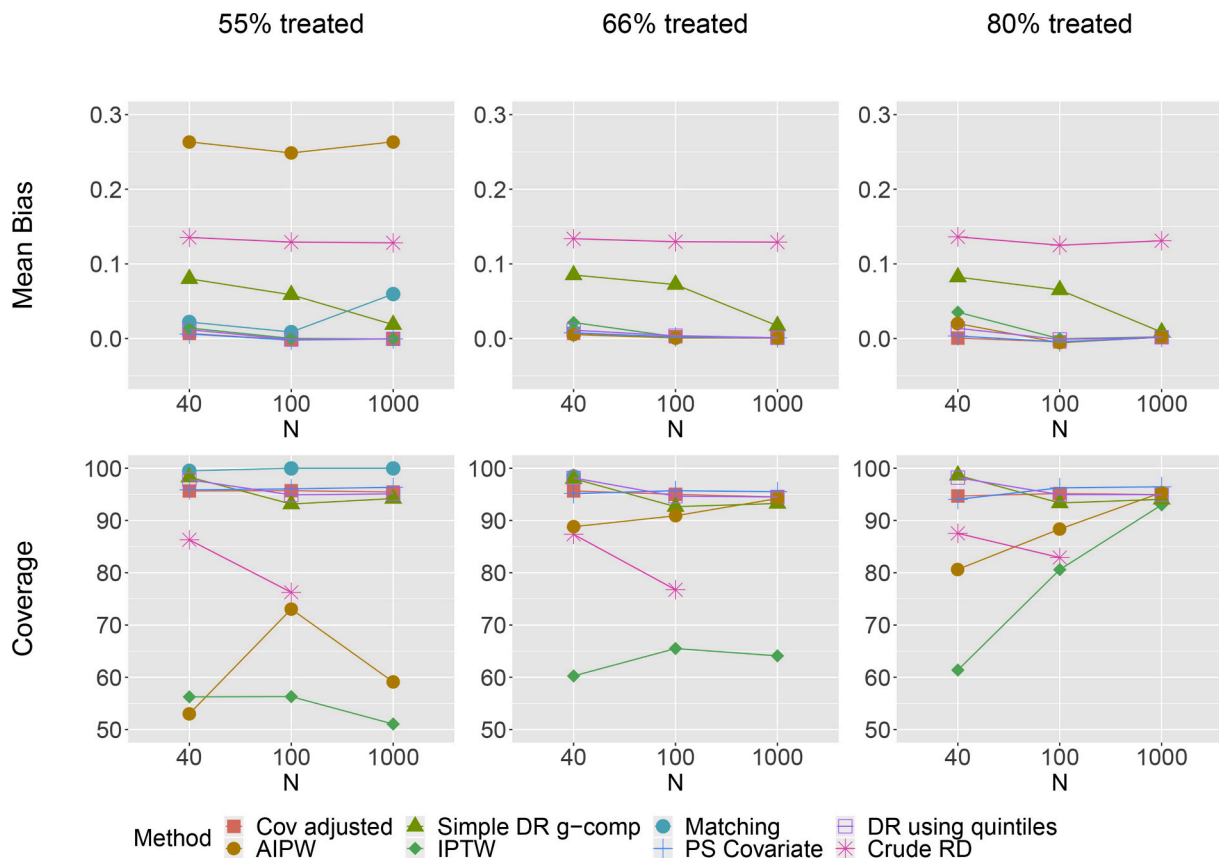


Fig. 3. Mean bias and coverage probabilities for Scenario 1 with a true RD of 0 and different proportions of treated individuals. Note that the coverage is truncated to $\geq 50\%$ implying that the unadjusted method is not displayed for $N = 1000$.

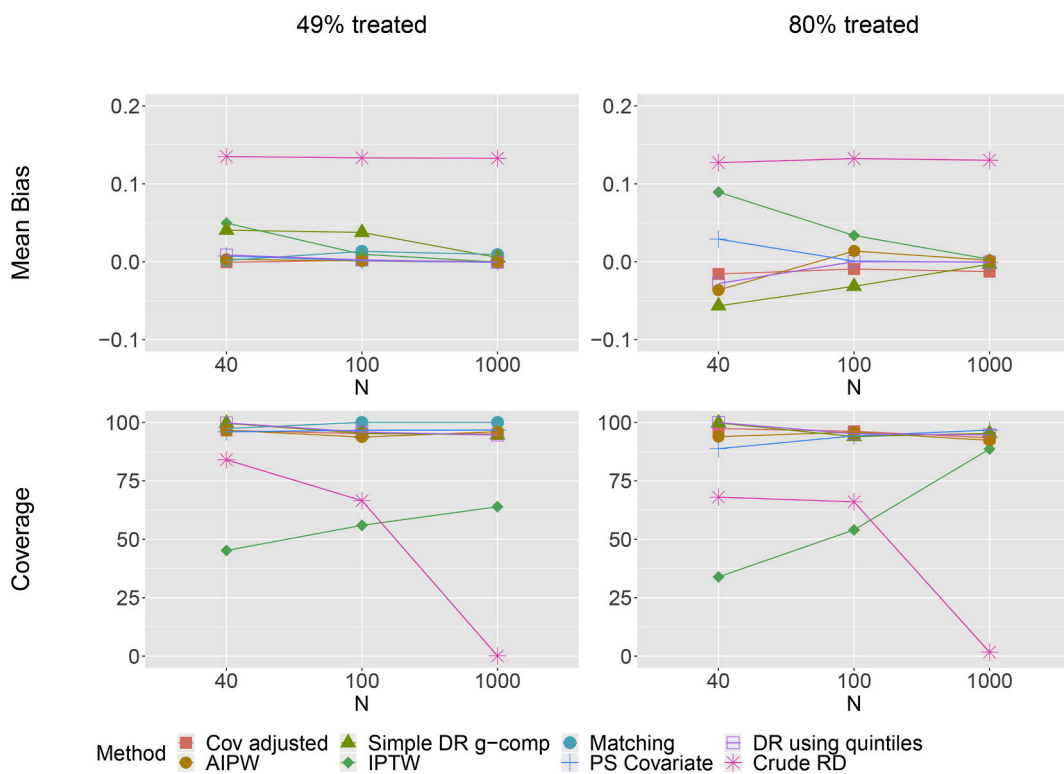


Fig. 4. Mean bias and coverage probabilities for Scenario 3 with a true RD of 0 and different proportions of treated individuals.

the results of the different PS methods.

Motivated by the study conducted by Gautret et al. [36] we investigated the properties of a range of causal inference methods in small samples. As expected this posed additional challenges to the various approaches. Interestingly, it turned out that the default settings in software implementations are often more suitable for large sample sizes and need to be adjusted for applications in small-scale studies. For example, we found that the matching procedure in R using the default calipers of 0 resulted in extremely biased results in our small sample simulations. SAS software, in contrast, uses a default caliper width of 0.25. The issue of choosing the right caliper width has recently been investigated by Wang [70], who recommended to take both matching and population bias into account.

We did not discuss the (causal) assumptions underlying the different estimation methods proposed in this paper. The recently published tutorial by Goetghebeur et al. [37] provides a general overview of these assumptions and how the methods discussed here invoke them. However, they also caution against the possible complications an applied statistician might face when conducting a causal analysis. In particular, it is not sufficient to focus on the non-randomized nature of a study and ignore, e. g. design issues, measurement error or study discontinuation, to name a few. This is exactly the case with our motivating data example. Focusing only on the non-randomized nature of the study by Gautret et al. [36], our reanalysis of the study (results not shown, code can be found on Github) was disappointing in that the conclusions based on various considered approaches did not differ from those reported by Gautret et al. [36] that we sought to correct and that disagree with recent large-scale trials [60]. Thus, our results demonstrate that while the causal inference methods can provide adjustment for baseline covariates, even a correctly applied causal inference method cannot compensate for design issues of the underlying study such as the small sample size, open label treatment and study discontinuations [58].

Our study has several limitations. First, while our simulation scenarios are carefully chosen to reflect different situations, we could only consider a limited number of settings. Thus, there is no guarantee that our results can be generalized to different situations. We therefore make our R code available, which can be used to explore specific scenarios. Second, as is known from the literature, we observed that the logistic regression model often failed for the small sample sizes, especially in combination with matching. To investigate whether PS matching can be improved in small sample sizes by using a penalization method shall be part of future research. Finally, our paper only studied a binary outcome and did not consider other commonly used outcomes in clinical studies such as time-to-event data or continuous outcomes. Especially in the context of time-to-event outcomes and longitudinal data, where time-varying treatments additionally complicate estimation, doubly robust g-computation such as TMLE [66,67] is recommended due to its good statistical properties. Similar to the issues discussed in this paper, the selection of a suitable endpoint requires some care and it is not always appropriate to use the most common approach, see [48] for a recent discussion in the context of time-to-event endpoints in COVID-19.

Besides the design of efficient trials to develop treatments for COVID-19 [64], one concern to trialists these days is the threat posed by the SARS-CoV-2 pandemic to clinical trials in non-COVID-19 indications [10,43]. SARS-CoV-2 infections of patients in these trials, or merely the increased risk thereof, might lead to post-randomization events (or intercurrent events in the language of the ICH E9 addendum [40]) such as treatment or study discontinuations as well as adverse events that ultimately invalidate an analysis relying on randomization. In such situations, the causal inference approach discussed here might provide a suitable alternative analysis strategy either as primary or sensitivity analysis.

Acknowledgement

Support by the DFG (grant FR 4121/2-1) is gratefully acknowledged.

Declaration of Competing Interest

The authors declare that they have no conflict of interest.

Appendix A. Supplementary material

Supplementary material to this article can be found online at <https://doi.org/10.1016/j.cct.2020.106213>.

References

- [1] Odd O. Aalen, Richard J. Cook, Kjetil Roysland, Does Cox analysis of a randomized survival study yield a causal treatment effect? *Lifetime Data Anal.* 21 (4) (2015) 579–593.
- [2] Alberto Abadie, Guido W. Imbens, Large sample properties of matching estimators for average treatment effects, *Econometrica* 74 (1) (2006) 235–267.
- [3] Alberto Abadie, Guido W. Imbens, On the failure of the bootstrap for matching estimators, *Econometrica* 76 (6) (2008) 1537–1557.
- [4] Alberto Abadie, Guido W. Imbens, Matching on the estimated propensity score, *Econometrica* 84 (2) (2016) 781–807.
- [5] Paul Elias Alexander, Victoria Borg Debono, Manoj J. Mammen, Alfonso Iorio, Komal Aryala, Dianna Deng, Eva Brocard, Waleed Alhazzani, COVID-19 coronavirus research has overall low methodological quality thus far: case in point for chloroquine/hydroxychloroquine, *J. Clin. Epidemiol.* 123 (2020) 120–126.
- [6] Robert P. Althaus, Donald Rubin, The computerized construction of a matched sample, *Am. J. Sociol.* 76 (2) (1970) 325–346.
- [7] Douglas G. Altman, *Statistics in Medical Journals: developments in the 1980s*, *Stat. Med.* 10 (12) (1991) 1897–1913.
- [8] Douglas G. Altman, Patrick Royston, What do we mean by validating a prognostic model? *Stat. Med.* 19 (4) (2000) 453–473.
- [9] Anais Andrillon, Romain Pirracchio, Sylvie Chevret, Performance of propensity score matching to estimate causal effects in small samples, *Stat. Methods Med. Res.* 29 (3) (2020) 644–658.
- [10] Stefan D Anker, Javed Butler, Muhammad Shahzeb Khan, William T Abraham, Johann Bauersachs, Edimar Bocchi, Biykem Bozkurt, Eugene Braunwald, Vijay K Chopra, John G Cleland, Justin Ezekowitz, Gerasimos Filippatos, Tim Friede, Adrian F Hernandez, Carolyn S P Lam, JoAnn Lindenfeld, John J V McMurray, Mandeep Mehra, Marco Metra, Milton Packer, Burkert Pieske, Stuart J Pocock, Piotr Ponikowski, Giuseppe M C Rosano, John R Teerlink, Hiroyuki Tsutsui, Dirk J Van Veldhuisen, Subodh Verma, Adriaan A Voors, Janet Wittes, Faiez Zannad, Jian Zhang, Petar Seferovic, and Andrew J S Coats. Conducting clinical trials in heart failure during (and after) the COVID-19 pandemic: an Expert Consensus Position Paper from the Heart Failure Association (HFA) of the European Society of Cardiology (ESC). *Eur. Heart J.* 41(22):2109–2117, 2020. URL: <https://doi.org/10.1093/eurheartj/ehaa461>, arXiv:<https://academic.oup.com/eurheartj/article-pdf/41/22/2109/33368354/ehaa461.pdf>, doi:<https://doi.org/10.1093/eurheartj/ehaa461>.
- [11] Peter C. Austin, The performance of different propensity score methods for estimating marginal odds ratios, *Stat. Med.* 26 (16) (2007) 3078–3094, <https://doi.org/10.1002/sim.2781>.
- [12] Peter C. Austin, The performance of different propensity-score methods for estimating differences in proportions (risk differences or absolute risk reductions) in observational studies, *Stat. Med.* 29 (20) (2010) 2137–2148.
- [13] Peter C. Austin, An introduction to propensity score methods for reducing the effects of confounding in observational studies, *Multivar. Behav. Res.* 46 (3) (2011) 399–424.
- [14] Peter C. Austin, A comparison of 12 algorithms for matching on the propensity score, *Stat. Med.* 33 (6) (2014) 1057–1069.
- [15] Peter C. Austin, Paul Grootendorst, Sharon-Lise T. Normand, Geoffrey M. Anderson, Conditioning on the propensity score can result in biased estimation of common measures of treatment effect: a Monte Carlo study, *Stat. Med.* 26 (4) (2007) 754–768, <https://doi.org/10.1002/sim.2618>.
- [16] Heejung Bang, James M. Robins, Doubly robust estimation in missing data and causal inference models, *Biometrics* 61 (4) (2005) 962–973.
- [17] Howard Bauchner, Phil B. Fontanarosa, Randomized Clinical Trials and COVID-19: Managing Expectations, *JAMA* 323 (22) (2020) 2262–2263, <https://doi.org/10.1001/jama.2020.8115>, arXiv:https://jamanetwork.com/journals/jama/articlepdf/2765696/jama_bauchner_2020_ed_200043.pdf.
- [18] Jan Beyersmann, Susanna Di Termini, Markus Pauly, Weak convergence of the wild bootstrap for the Aalen–Johansen estimator of the cumulative incidence function of a competing risk, *Scand. J. Stat.* 40 (3) (2013) 387–402.
- [19] J. Martin Bland, Douglas G. Altman, The odds ratio, *BMJ* 320 (7247) (2000) 1468.
- [20] Tobias Bluhmki, Claudia Schmoor, Dennis Dobler, Markus Pauly, Jürgen Finke, Martin Schumacher, Jan Beyersmann, A wild bootstrap approach for the Aalen–Johansen estimator, *Biometrics* 74 (3) (2018) 977–985, <https://doi.org/10.1111/biom.12861>.
- [21] David R. Boulware, Matthew F. Pullen, Ananta S. Bangdiwala, Katelyn A. Pastick, Sarah M. Lofgren, Elizabeth C. Okafor, Caleb P. Skipper, Alanna A. Nascene, Melanie R. Nicol, Mahsa Abassi, Nicole W. Engen, Matthew P. Cheng, Derek LaBar, Sylvain A. Lother, Lauren J. MacKenzie, Glen Drobot, Nicole Marten, Ryan Zarychanski, Lauren E. Kelly, Ilan S. Schwartz, Emily G. McDonald, Radha Rajasingham, Todd C. Lee, and Kathy H. Hullsiek. A randomized trial of hydroxychloroquine as postexposure prophylaxis for covid-19. *New England J. Med.*, 2020. arXiv:<https://doi.org/10.1056/NEJMoa2016638>, doi:<https://doi.org/10.1056/NEJMoa2016638>.

- [22] M. Alan Brookhart, Sebastian Schneeweiss, Kenneth J. Rothman, Robert J. Glynn, Jerry Avorn, Til Stürmer, Variable selection for propensity score models, *Am. J. Epidemiol.* 163 (12) (2006) 1149–1156.
- [23] Alexandre B. Cavalcanti, Fernando G. Zampieri, Regis G. Rosa, Luciano C.P. Azevedo, Viviane C. Veiga, Alvaro Avezum, Lucas P. Damiani, Aline Marcadenti, Leticia Kawano-Dourado, Thiago Lisboa, Debora L. M. Junqueira, de Barros Silva e Pedro G.M., Lucas Tramuja, Erlon O. Abreu-Silva, Ligia N. Laranjeira, Aline T. Soares, Leandro S. Echenique, Adriano J. Pereira, Flavio G.R. Freitas, Otavio C.E. Gebara, Vicente C.S. Dantas, Remo H.M. Furtado, Eveline P. Milan, Nicole A. Golin, Fabio F. Cardoso, Israel S. Maia, Conrado R. Hoffmann Filho, Adrian P.M. Kormann, Roberto B. Amazonas, Monalisa F. Bocchi de Oliveira, Ary Serpa-Neto, Maicon Falavigna, Renato D. Lopes, Flavia R. Machado, and Otavio Berwanger. Hydroxychloroquine with or without azithromycin in mild-to-moderate covid-19. *New England J. Med.*, 2020. arXiv: 10.1056/NEJMoa2019014.
- [24] Yin Bun Cheung, A modified least-squares regression approach to the estimation of risk difference, *Am. J. Epidemiol.* 166 (11) (2007) 1337–1344.
- [25] Thomas D. Cook, Advanced statistics: up with odds ratios! A case for odds ratios when outcomes are common, *Acad. Emerg. Med.* 9 (12) (2002) 1430–1434.
- [26] Andrea Cortegiani, Giulia Ingoglia, Mariachiara Ippolito, Antonino Giarratano, Sharon Einav, A systematic review on the efficacy and safety of chloroquine for the treatment of COVID-19, *J. Crit. Care* 57 (2020) 279–283.
- [27] Peter Cummings, The relative merits of risk ratios and odds ratios, *Arch. Pediatr. Adolesc. Med.* 163 (5) (2009) 438–445.
- [28] Awa Dieng, Yameng Liu, Sudeepa Roy, Cynthia Rudin, Alexander Volfvsky, Interpretable almost-exact matching for causal inference, *Proceed. Machine Learn. Res.* 89 (2019) 2445.
- [29] Bradley Efron, Robert Tibshirani, Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy, *Stat. Sci.* (1986) 54–75.
- [30] Matthew E. Falagas, Gregory C. Makris, Drosos E. Karageorgopoulos, Maria Batsiou, Vangelis G. Alexiou, How well do clinical researchers understand risk estimates? *Epidemiology* 20 (6) (2009) 930–931.
- [31] João Pedro Ferreira, Murray Epstein, Faiez Zannad, The decline of the experimental paradigm during the COVID-19 pandemic: a template for the future, *Am. J. Med.* (2020), <https://doi.org/10.1016/j.amjmed.2020.08.021>.
- [32] Sarah Friedrich, Frank Konietzschke, Markus Pauly, A wild bootstrap approach for nonparametric repeated measurements, *Comput. Stat. Data Anal.* 113 (2017) 38–52.
- [33] Sarah Friedrich, Markus Pauly, MATS: inference for potentially singular and heteroscedastic MANOVA, *J. Multivar. Anal.* 165 (2018) 166–179.
- [34] Christian Funck-Brentano, Lee S. Nguyen, Joe-Elie Salem, Retraction and republication: cardiac toxicity of hydroxychloroquine in COVID-19, *The Lancet* (2020), [https://doi.org/10.1016/S0140-6736\(20\)31528-2](https://doi.org/10.1016/S0140-6736(20)31528-2).
- [35] Mitchell H. Gail, S. Wieand, Steven Piantadosi, Biased estimates of treatment effect in randomized experiments with nonlinear regressions and omitted covariates, *Biometrika* 71 (3) (1984) 431–444.
- [36] Philippe Gautre, Jean-Christophe Lagier, Philippe Parola, Line Meddeb, Morgane Mailhe, Barbara Doudier, Johan Courjon, Valérie Giordanengo, Vera Esteves Vieira, Hervé Tisot Dupont, et al., Hydroxychloroquine and azithromycin as a treatment of COVID-19: results of an open-label non-randomized clinical trial, *Int. J. Antimicrob. Agents* 56 (1) (2020), <https://doi.org/10.1016/j.ijantimicag.2020.105949>.
- [37] Els Goetghebeur, Saskia le Cessie, Bianca De Stavola, Erica EM Moodie, Ingeborg Waernbaum, and “on behalf of” the topic group Causal Inference (TG7) of the STRATOS initiative. Formulating causal questions and principled statistical answers. *Stat. Med.*, pages 1–27, 2020. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.8741>, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/sim.8741>, doi:<https://doi.org/10.1002/sim.8741>.
- [38] Miguel Hernan, James Robins, *Causal Inference: What if*, Chapman & Hall/CRC, Boca Raton, 2020.
- [39] Peter Horby, Marion Mafham, Louise Linsell, Jennifer L Bell, Natalie Staplin, Jonathan R Emberson, Martin Wiselka, Andrew Ustianowski, Einas Elmahi, Benjamin Prudon, Anthony Whitehouse, Timothy Felton, John Williams, Jakki Faccenda, Jonathan Underwood, J Kenneth Baillie, Lucy Chappell, Saul N Faust, Thomas Jaki, Katie Jeffery, Wei Shen Lim, Alan Montgomery, Kathryn Rowan, Joel Tarning, James A Watson, Nicholas J White, Edmund Juszczak, Richard Haynes, and Martin J Landray. Effect of Hydroxychloroquine in Hospitalized Patients with COVID-19: Preliminary results from a multi-centre, randomized, controlled trial. medRxiv, 2020. URL: <https://www.medrxiv.org/content/early/2020/07/15/2020.07.15.20151852>, arXiv:<https://www.medrxiv.org/content/early/2020/07/15/2020.07.15.20151852.full.pdf>, doi:<https://doi.org/10.1101/2020.07.15.20151852>.
- [40] ICH. ICH E9 (R1) addendum on estimands and sensitivity analysis in clinical trials to the guideline on statistical principles for clinical trials. 2019. URL: <https://www.ema.europa.eu/en/ich-e9-statistical-principles-clinical-trials>.
- [41] Joseph D.Y. Kang, Joseph L. Schafer, Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data, *Stat. Sci.* 22 (4) (2007) 523–539.
- [42] Frank Konietzschke, Arne C. Bathke, Solomon W. Harrar, Markus Pauly, Parametric and nonparametric bootstrap methods for general MANOVA, *J. Multivar. Anal.* 140 (2015) 291–301.
- [43] Cornelia Ursula Kunz, Silke Jörgens, Frank Bretz, Nigel Stallard, Kelly Van Lancker, Dong Xi, Sarah Zohar, Christoph Gerlinger, and Tim Friede. Clinical trials impacted by the covid-19 pandemic: adaptive designs to the rescue? *Stat. Biopharmaceut. Res.*, 2020. URL: <https://doi.org/10.1080/19466315.2020.1799857>, arXiv:<https://doi.org/10.1080/19466315.2020.1799857>.
- [44] Huzhang Mao, Li Liang, Flexible regression approach to propensity score analysis and its relationship with matching and weighting, *Stat. Med.* (2020), <https://doi.org/10.1002/sim.8526>.
- [45] Huzhang Mao, Li Liang, Tom Greene, Propensity score weighting analysis and treatment effect discovery, *Stat. Methods Med. Res.* 28 (8) (2019) 2439–2454.
- [46] Torben Martinussen, Stijn Vansteelandt, On collapsibility and confounding bias in Cox and Aalen regression models, *Lifetime Data Anal.* 19 (3) (2013) 279–296.
- [47] Daniel F. McCaffrey, Beth Ann Griffin, Daniel Almirall, Mary Ellen Slaughter, Rajeev Ramchand, Lane F. Burgette, A tutorial on propensity score estimation for multiple treatments using generalized boosted models, *Stat. Med.* 32 (19) (2013) 3388–3414.
- [48] Zachary R. McCaw, Lu Tian, Kevin N. Sheth, Wan-Ting Hsu, W. Taylor Kimberley, Lee-Jen Wei, Selecting appropriate endpoints for assessing treatment effects in comparative clinical studies for COVID-19, *Contemp. Clin. Trials* 97 (2020) 106145, <https://doi.org/10.1016/j.cct.2020.106145>.
- [49] Tobias Mütze, Tim Friede, Data monitoring committees for clinical trials evaluating treatments of COVID-19, *Contemp. Clin. Trials* 98 (2020) 106154, <https://doi.org/10.1016/j.cct.2020.106154>.
- [50] Markus Pauly, Edgar Brunner, Frank Konietzschke, Asymptotic permutation tests in general factorial designs, *J. R. Stat. Soc. Ser. B (Statistical Methodology)* 77 (2) (2015) 461–473.
- [51] Romain Pirracchio, Matthieu Resche-Rigon, Sylvie Chevret, Evaluation of the propensity score methods for estimating marginal odds ratios in case of small sample size, *BMC Med. Res. Methodol.* 12 (1) (2012) 70.
- [52] Rainer Puhrt, Georg Heinze, Mariana Nold, Lara Lusa, Angelika Geroldinger, Firth’s logistic regression with rare events: accurate effect estimates and predictions? *Stat. Med.* 36 (14) (2017) 2302–2317.
- [53] James Robins, A new approach to causal inference in mortality studies with a sustained exposure period – application to control of the healthy worker survivor effect, *Math. Modell.* 7 (9–12) (1986) 1393–1512.
- [54] Laurence D. Robinson, Nicholas P. Jewell, Some surprising results about covariate adjustment in logistic regression models, *Int. Stat. Rev. Rev. Int. Stat.* (1991) 227–240.
- [55] Bejammn N. Rome, Jerry Avorn, Drug evaluation during the covid-19 pandemic, *N. Engl. J. Med.* 382 (24) (2020) 2282–2284.
- [56] Paul R. Rosenbaum, Donald B. Rubin, The central role of the propensity score in observational studies for causal effects, *Biometrika* 70 (1) (1983) 41–55.
- [57] Paul R. Rosenbaum, Donald B. Rubin, Reducing bias in observational studies using subclassification on the propensity score, *J. Am. Stat. Assoc.* 79 (387) (1984) 516–524.
- [58] Donald B. Rubin, For objective causal inference, design trumps analysis, *Ann. Appl. Stat.* 2 (3) (2008) 808–840.
- [59] Donald B. Rubin, Neal Thomas, Matching using estimated propensity scores: relating theory to practice, *Biometrics* (1996) 249–264.
- [60] Sebastian E. Sattui, Jean W. Liew, Elizabeth R. Graef, Ariella Coler-Reilly, Francis Berenbaum, Ali Duarte-Garcia, Carly Harrison, Maximilian F. Konig, Peter Korsten, Michael S. Putman, Philip C. Robinson, Emily Siroitch, Manuel F. Ugarte-Gil, Kate Webb, Kristen J. Young, Alfred H.J. Kim, Jeffrey A. Sparks, Swinging the pendulum: lessons learned from public discourse concerning hydroxychloroquine and COVID-19, *Exp. Rev. Clin. Immunol.* (2020), <https://doi.org/10.1080/17446666X.2020.1792778>.
- [61] Arvid Sjölander, Elisabeth Dahlqvist, Johan Zetterqvist, A note on the noncollapsibility of rate differences and rate ratios, *Epidemiology* 27 (3) (2016) 356–359.
- [62] Jonathan M. Snowden, Sherri Rose, Kathleen M. Mortimer, Implementation of G-computation on a simulated data set: demonstration of a causal inference technique, *Am. J. Epidemiol.* 173 (7) (2011) 731–738.
- [63] Jeffrey Sonis, Odds ratios vs risk ratios, *JAMA* 320 (19) (2018) 2041, <https://doi.org/10.1001/jama.2018.14417>, arXiv:https://jamanetwork.com/journals/jama/articlepdf/2715584/jama_sonis_2018_le_180130.pdf.
- [64] Nigel Stallard, Lisa Hampson, Norbert Benda, Werner Brannath, Thomas Burnett, Tim Friede, Peter K. Kimani, Franz Koenig, Johannes Krisam, Pavel Mozgunov, Martin Posch, James Wason, Gernot Wassmer, John Whitehead, S. Faye Williamson, Sarah Zohar, Thomas Jaki, Efficient adaptive designs for clinical trials of interventions for COVID-19, *Stat. Biopharmaceut. Res.* 12 (4) (2020) 483–497, <https://doi.org/10.1080/19466315.2020.1790415>.
- [65] Erout W. Steyerberg, *Clinical Prediction Models*, Springer, 2019.
- [66] Ori M. Stitelman, Victor De Gruttola, Mark J. van der Laan, A general implementation of tmle for longitudinal data applied to causal inference in survival analysis, *Int. J. Biostat.* 8 (1) (2012).
- [67] Mark J. van der Laan, Susan Gruber, Collaborative double robust targeted maximum likelihood estimation, *Int. J. Biostat.* 6 (1) (2010).
- [68] Maarten van Smeden, Joris A.H. de Groot, Karel G.M. Moons, Gary S. Collins, Douglas G. Altman, Marinus J.C. Eijkemans, Johannes B. Reitsma, No rationale for 1 variable per 10 events criterion for binary logistic regression analysis, *BMC Med. Res. Methodol.* 16 (1) (2016) 163.
- [69] Maarten van Smeden, Karel G.M. Moons, Joris A.H. de Groot, Gary S. Collins, Douglas G. Altman, Marinus J.C. Eijkemans, Johannes B. Reitsma, Sample size for binary logistic prediction models: beyond events per variable criteria, *Stat. Methods Med. Res.* 28 (8) (2019) 2455–2474.
- [70] Jixian Wang, To use or not to use propensity score matching? *Pharm. Stat.* (2020) <https://doi.org/10.1002/pst.2051>.
- [71] Tianyu Wang, Marco Morucci, M. Awan, Yameng Liu, Sudeepa Roy, Cynthia Rudin, Alexander Volfvsky, Flame: A Fast Large-scale almost Matching Exactly Approach to Causal Inference, arXiv preprint, 2017. arXiv:1707.06315.
- [72] Sherry Weitzen, Kate L. Lapane, Alicia Y. Toledano, Anne L. Hume, Vincent Mor, Weaknesses of goodness-of-fit tests for evaluating propensity score models: the case of the omitted confounder, *Pharmacoepidemiol. Drug Saf.* 14 (4) (2004) 227–238, 10.1002%2Fpds.986.