HHS Public Access

Author manuscript

Epidemiology. Author manuscript; available in PMC 2021 September 01.

Published in final edited form as:

Epidemiology. 2020 September ; 31(5): 692–694. doi:10.1097/EDE.0000000000001234.

Defining and Identifying Per-Protocol Effects in Randomized Trials

Jacqueline E. Rudolph[1], **Ashley I. Naimi**[1], **Daniel J. Westreich**[2], **Edward H. Kennedy**[3], **Enrique F. Schisterman**[4]

[1]Department of Epidemiology, University of Pittsburgh.

[2]Department of Epidemiology, University of North Carolina at Chapel Hill.

[3]Department of Statistics, Carnegie Mellon University.

[4]Epidemiology Branch, Division of Intramural Population Health Research, *Eunice Kennedy Shriver* National Institute of Child Health and Human Development, Bethesda, MD, United States

Abstract

In trials with non-compliance to assigned treatment, researchers might be interested in estimating a per-protocol effect – a comparison of two counterfactual outcomes defined by treatment assignment and (often time-varying) compliance with a well-defined treatment protocol. Here, we provide a general counterfactual definition of a per-protocol effect and discuss examples of per-protocol effects that are of either substantive or methodologic interest. In doing so, we seek to make more concrete what per-protocol effects are and highlight that one can estimate per-protocol effects that are more than just a comparison of always taking treatment in two distinct treatment arms. We then discuss one set of identifiability conditions that allow for identification of a causal per-protocol effect, highlighting some potential violations of those conditions that might arise when estimating per-protocol effects.

Keywords

per-protocol effect; randomized controlled trials; noncompliance; causal inference; identifiability

Introduction

In randomized controlled trials (RCTs) with non-compliance, the intention-to-treat effect cannot be interpreted as an estimate of treatment efficacy.[1] In such cases, researchers may wish to estimate effects that would have been observed if study participants had (possibly contrary to fact) followed a pre-designated protocol. These per-protocol effects require careful consideration and handling of the potentially time-varying causes of non-compliance, as has been discussed in previous papers.[1–5]

Correspondence: Jacqueline E. Rudolph, Department of Epidemiology, Graduate School of Public Health, University of Pittsburgh, 130 DeSoto Street, Pittsburgh, PA 15261, 412-624-3056, jer186@pitt.edu.

Conflicts of Interest: None declared

However, prior work on per-protocol effects has rarely (if ever) explicitly stated the counterfactual definition of the target parameter, which could lead to some confusion regarding the definition of a per-protocol effect. There are also different target per-protocol parameters that can be used to answer different research questions or provide insight regarding the validity of our models or causal assumptions. Thus, we here provide a general counterfactual definition and describe several example per-protocol effects. We then comment on one sufficient set of conditions needed to identify these effects, as these have also not been explicitly discussed in relation to per-protocol effects.

Defining per-protocol effects

Let R denote randomized treatment. Participants are followed until outcome Y or end of follow-up. Compliance is assessed in time intervals ($J = 1,\ldots,m$) where $C_j = 1$ means the participant was compliant for time-point j. Looking at compliance across follow-up, we also define compliance history $\bar{C}$ for each participant. While our example is general, note that estimating per-protocol effects requires clearly defined compliance protocols that directly correlate with the research question of interest and with substantive knowledge. Other work has explored in detail best practices for defining protocols.[5]

We define a per-protocol effect as:

$$E\left[Y^{r=p_1,\bar{c}=\gamma_1}\right] - E\left[Y^{r=p_2,\bar{c}=\gamma_2}\right]$$

Where $r = p$ is a given randomized treatment and $\bar{c} = \gamma$ is a given compliance history. Suppose we were interested in two trial arms, $R=\{0,1\}$ and $\bar{C}$ defined by always- or never-complying with the same protocol regardless of arm. We can then define six basic per-protocol effects (Table), three of which we discuss here.

Under relevant identifiability conditions (listed below), contrast A is the effect of treatment relative to comparator under full compliance to the protocol and will most closely reflect the effect of the treatment's active ingredient. If the protocol is static and requires taking treatment every day, A would be the maximum possible effect. If a dynamic protocol is specified (e.g., take treatment until some clinical event occurs), this effect will best approximate the biologic effect over the relevant time scale. A is arguably the effect of primary interest in many settings and has been the target parameter in most modern per-protocol analyses.[1,6]

It is additionally possible to use the other parameters to triangulate this effect.[7,8] Subtracting C from B yields A due to cancellation of the $E\left[Y^{r=0,\bar{c}=0}\right]$ term, as such:

$$\left\{E\left[Y^{r=1,\bar{c}=1}\right] - E\left[Y^{r=0,\bar{c}=0}\right]\right\} - \left\{E\left[Y^{r=0,\bar{c}=1}\right] - E\left[Y^{r=0,\bar{c}=0}\right]\right\} = E\left[Y^{r=1,\bar{c}=1}\right] - E\left[Y^{r=0,\bar{c}=1}\right]$$

$$B - C = A$$

Similar triangulation can be obtained with D and E. The benefit of triangulating A in this way is methodologic. If the results from taking B – C and D – E were reasonably similar to A (acknowledging that random error might make them non-identical), we might feel reassurance that we had specified our statistical or causal models correctly. However, if estimates differ, this signals potential problems that should be explored further, although we cannot ascertain the exact source of error.

Contrast B is the joint effect of being assigned to treatment and always-complying with the protocol, relative to assignment to the comparator and never-complying. This effect will capture a combination of the treatment's biologic effect and the effect of protocol compliance. B could be used to assess the maximum benefit under various protocols, perhaps by varying the protocol of the treatment from "take every day for all of follow-up" to "take every day for a specific time period" to "take x times per week." Unlike A, this effect includes the impact of being engaged with a treatment protocol, which might be valuable to understand before applying a protocol in a real-world setting. Comparing B to A could reveal whether the act of staying compliant with a given protocol leads to additional benefit (or harm) beyond that seen from the biologic effect alone.

F is the effect of being assigned to treatment and never-complying with the treatment protocol versus being assigned to comparator and never-complying with the comparator protocol and is useful as a tool for model validation. In a study where one expects any sub-optimal compliance to mean the participant will experience no effect of treatment or comparator, F should be null (or close to null, given the possibility of random error). While previous studies have used contrast C for model validation,[9,10] a non-null F could also indicate problems with the statistical or causal models that may be undermining the validity of other effects estimates.

For information on how to estimate per-protocol effects, we point readers toward several published papers, including Lodi, Cain, Murray, and Toh.[6,9,11,12] While these papers have focused on estimating contrast A, the methods could be used to estimate any of the contrasts.

Identifying per-protocol effects

One sufficient set of conditions to identify a causal effect is exchangeability, positivity, and counterfactual consistency.

Exchangeability (specifically, conditional exchangeability) requires the potential outcomes $Y^{r,\bar{c}}$ to be independent of observed treatment and compliance at each time-point, conditional on compliance history prior to that time point and a set of covariates L representing confounder history.[13–16] Any variables that affect compliance and the outcome should be included in L. Particular consideration should be given to whether L will differ by treatment arm or protocol.

Positivity requires the probability of being exposed or unexposed, given *L*, be bounded away from 0 and one; we must here consider exposure defined by randomization and compliance. If the per-protocol effect of interest was the comparison of a given treatment under always-complying versus never-complying, it is possible that there would be very few individuals who were assigned to that treatment who genuinely never took the treatment according to the rule. Another challenge to positivity could arise because a large number of participants in a given arm choose to no longer comply with the treatment because of side effects (typically included in *L*). This is one reason to specify a protocol in which a participant is still compliant if they stopped treatment after experiencing a side effect. However, this decision should be dictated by substantive concerns – not merely because positivity is violated.

Counterfactual consistency requires that any variation in how participants did or did not comply with the protocol is irrelevant for the effect on the outcome.[17] Such variations could include taking treatment 7 days instead of 5 days or taking treatment with water, fruit juice, or a caffeinated beverage. Consistency might be particularly difficult to assume when the comparator group is "standard of care," which could be defined differently from patient to patient.

Beyond the above, we typically also assume no model misspecification. In RCTs with noncompliance, data could be complex, due to the presence of numerous time points and a large set *L*. In such cases, researchers often use parametric models, which require strict, often unrealistic model form assumptions. Indeed, the primary methods that have been used to estimate per-protocol effects, namely inverse probability weighting and g-computation, [1,6,11,15,18–21] generally require the use of parametric models. Unlike the above, though, violations of correct model specification are a statistical, rather than causal, concern.[22]

Any of these assumptions may not be credible for a given trial, and there are a number of practical limitations which could preclude one from estimating or identifying the described per-protocol effects.[5] For instance, insufficient samples of continuously non-adherent participants could make estimating parameter F impractical. Small sample sizes in general could mean one would be underpowered to detect any difference between parameter A or one of its triangulation parameters – not to mention the difficulty in assuming positivity. Additionally, one might lack the data one would need to identify the effect of interest. In particular, a trial might not measure all the confounders of the relationship between adherence and the outcome, and adherence itself may be poorly measured. Such practical limitations must be considered prior to estimating any per-protocol effect.

Regardless, these conditions (or a different sufficient set)[23] are necessary to identify per-protocol effects. While most of the assumptions are not testable using observed data, sensitivity analyses can be conducted,[24,25] and there exist estimators that can relax some of these assumptions. For example, the positivity assumption is not required for incremental propensity score effects,[26] while machine learning can be used with double robust estimators to relax parametric modeling assumptions.[8,27–30]

Discussion

When assessing the role of compliance to assigned treatment in an RCT, researchers can define a wide range of protocols and per-protocol effects. We here provided potential contrasts that might be of substantive or methodologic interest and discussed several important considerations when attempting to identify per-protocol effects. The key piece underlying per-protocol effect estimation is the specification of the protocol, which should primarily be based on the research question. All other considerations, including the contrast(s) of interest, estimation approach, and ability to meet identifiability conditions, follow from the chosen protocol.

Sources of Funding:

This work was supported by grant NIH-R01 HD093602 and the Intramural Research Program of the *Eunice Kennedy Shriver* National Institute of Child Health and Human Development, National Institutes of Health, Bethesda, Maryland (contract numbers HHSN267200603423, HHSN267200603424, and HHSN267200603426).

References

1. Hernán MA, Robins JM. Per-Protocol Analyses of Pragmatic Trials. N Engl J Med 2017;377(14):1391–1398. [PubMed: 28976864]
2. Hernán MA, Hernandez-Diaz S. Beyond the intention-to-treat in comparative effectiveness research. Clin Trials 2012;9(1):48–55. [PubMed: 21948059]
3. Shrier I, Steele RJ, Verhagen E, Herbert R, Riddell CA, Kaufman JS. Beyond intention to treat: what is the right question? Clin Trials. 2014;11(1):28–37. [PubMed: 24096636]
4. Dodd S, White IR, Williamson P. Nonadherence to treatment protocol in published randomised controlled trials: a review. Trials. 2012;13:84. [PubMed: 22709676]
5. Murray EJ, Swanson SA, Hernán MA. Guidelines for estimating causal effects in pragmatic randomized trials. https://arxiv.org/pdf/1911.06030.pdf. Published 2019. Accessed2020.
6. Lodi S, Sharma S, Lundgren JD, et al. The per-protocol effect of immediate versus deferred antiretroviral therapy initiation. AIDS. 2016;30(17):2659–2663. [PubMed: 27782964]
7. Lawlor DA, Tilling K, Davey Smith G. Triangulation in aetiological epidemiology. Int J Epidemiol. 2016;45(6):1866–1886. [PubMed: 28108528]
8. Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. Biometrics. 2005;61(4):962–973. [PubMed: 16401269]
9. Murray EJ, Hernán MA. Adherence adjustment in the Coronary Drug Project: A call for better per-protocol effect estimates in randomized trials. Clin Trials. 2016;13(4):372–378. [PubMed: 26951361]
10. Murray EJ, Hernán MA. Improved adherence adjustment in the Coronary Drug Project. Trials. 2018;19(1):158. [PubMed: 29506561]
11. Cain LE, Cole SR. Inverse probability-of-censoring weights for the correction of time-varying noncompliance in the effect of randomized highly active antiretroviral therapy on incident AIDS or death. Stat Med. 2009;28(12):1725–1738. [PubMed: 19347843]
12. Toh S, Hernandez-Diaz S, Logan R, Robins JM, Hernán MA. Estimating absolute risks in the presence of nonadherence: an application to a follow-up study with baseline randomization. Epidemiology. 2010;21(4):528–539. [PubMed: 20526200]
13. Robins JM, Hernán MA. Estimation of the causal effects of time-varying exposures In: Fitzmaurice G, Davidian M, Verbeke G, Molenberghs G, eds. Advances in Longitudinal Data Analysis. New York: Chapman and Hall/CRC Press; 2009.
14. Naimi AI, Cole SR, Kennedy EH. An introduction to g methods. Int J Epidemiol. 2017;46(2):756–762. [PubMed: 28039382]

15. Hernán MA, Hernandez-Diaz S, Robins JM. Randomized trials analyzed as observational studies. Ann Intern Med. 2013;159(8):560–562. [PubMed: 24018844]
16. Daniel RM, Cousens SN, De Stavola BL, Kenward MG, Sterne JA. Methods for dealing with time-dependent confounding. Stat Med. 2013;32(9):1584–1618. [PubMed: 23208861]
17. VanderWeele TJ. Concerning the consistency assumption in causal inference. Epidemiology. 2009;20(6):880–883. [PubMed: 19829187]
18. Robins J A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. Mathematical Modelling. 1986;7(9):1393–1512.
19. Robins JM, Hernán MA, Brumback B. Marginal structural models and causal inference in epidemiology. Epidemiology. 2000;11(5):550–560. [PubMed: 10955408]
20. Hernán MA, Alonso A, Logan R, et al. Observational studies analyzed like randomized experiments: an application to postmenopausal hormone therapy and coronary heart disease. Epidemiology. 2008;19(6):766–779. [PubMed: 18854702]
21. Naimi AI, Perkins NJ, Mumford SL, et al. The per protocol effect of preconception-initiated low-dose aspirin on hCG pregnancy, pregnancy loss, and live birth: a randomized trial. In:2019.
22. Kennedy EH. Nonparametric efficiency theory and machine learning in causal inference. http://www.ehkennedy.com/uploads/5/8/4/5/58450265/tutorial.pdf. Published 2018. Accessed2019.
23. Kennedy EH. Semiparametric theory and empirical processes in causal inference. https://arxiv.org/abs/1510.04740.Published 2016 Updated July 22, 2016 Accessed2020.
24. Vanderweele TJ, Arah OA. Bias formulas for sensitivity analysis of unmeasured confounding for general outcomes, treatments, and confounders. Epidemiology. 2011;22(1):42–52. [PubMed: 21052008]
25. Robins JM, Rotnitzky A, Scharfstein DO. Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models In: Halloran ME, Berry D, eds. Statistical Models in Epidemiology, the Environment, and Clinical Trials. New York: Springer; 2000:1–94.
26. Kennedy EH. Nonparametric Causal Effects Based on Incremental Propensity Score Interventions. Journal of the American Statistical Association. 2019;114(526):645–656.
27. Kennedy EH. Semiparametric theory. https://arxiv.org/abs/1709.06418.Published 2017 Accessed2019.
28. Naimi AI, Kennedy EH. Nonparametric double robustness. https://arxiv.org/abs/1711.07137.Published 2017 Accessed2019.
29. Tsiatis AA. Semiparametric Theory and Missing Data. New York: Springer-Verlag; 2006.
30. van der Laan MJ, Rose S. Targeted Learning in Data Science. In: Springer International Publishing AG; 2018.

Table.

Summary of six basic per-protocol parameters given randomization to a binary treatment and constant compliance or noncompliance across follow-up

Parameters	Counterfactual Notation
A	$E\left[Y^{r}=1,\bar{c}=1\right]-E\left[Y^{r}=0,\bar{c}=1\right]$
B	$E\left[Y^{r}=1,\bar{c}=1\right]-E\left[Y^{r}=0,\bar{c}=0\right]$
C	$E\left[Y^{r}=0,\bar{c}=1\right]-E\left[Y^{r}=0,\bar{c}=0\right]$
D	$E\left[Y^{r}=1,\bar{c}=1\right]-E\left[Y^{r}=1,\bar{c}=0\right]$
E	$E\left[Y^{r}=0,\bar{c}=1\right]-E\left[Y^{r}=1,\bar{c}=0\right]$
F	$E\left[Y^{r}=1,\bar{c}=0\right]-E\left[Y^{r}=0,\bar{c}=0\right]$