Published in final edited form as:

Epidemiology. 2017 January ; 28(1): 54–59. doi:10.1097/EDE.0000000000000564.

Biases in randomized trials: a conversation between trialists and epidemiologists

Mohammad Ali Mansournia¹, Julian P. T. Higgins², Jonathan A. C. Sterne², and Miguel A. Hernán^{3,4}

¹Department of Epidemiology and Biostatistics, School of Public Health, Tehran University of Medical Sciences, Tehran, Iran

²School of Social and Community Medicine, University of Bristol, Bristol, UK

³Departments of Epidemiology and Biostatistics, Harvard School of Public Health, Boston, MA, USA

⁴Harvard-MIT Division of Health Sciences and Technology, Boston, MA, USA

Abstract

Trialists and epidemiologists often employ different terminology to refer to biases in randomized trials and observational studies, even though many biases have a similar structure in both types of study. We use causal diagrams to represent the structure of biases, as described by the Cochrane Collaboration for randomized trials, and provide a translation to the usual epidemiologic terms of confounding, selection bias, and measurement bias. This structural approach clarifies that an explicit description of the inferential goal—the intention-to-treat effect or the per-protocol effect—is necessary to assess risk of bias in the estimates. Being aware of each other's terminologies will enhance communication between trialists and epidemiologists when considering key concepts and methods for causal inference.

Keywords

bias; randomized trial; Cochrane bias domains; causal diagram

Introduction

Randomized controlled trials (RCTs) and observational studies are used to assess the causal effects of medical interventions.¹ By definition, treatment strategies are randomly assigned in RCTs but not in observational studies. Randomization, which prevents bias due to non-comparability between groups, is exploited in full when the data analysis follows the “intention-to-treat” principle.

Address for correspondence: Mohammad Ali Mansournia, Department of Epidemiology and Biostatistics, School of Public Health, Tehran University of Medical Sciences, PO Box: 14155-6446, Tehran, Iran. Tel: +98-21-88989123; Fax: +98-21-88989127; mansournia_ma@yahoo.com.

Conflicts of Interest: The authors declare no conflict of interest.

Another difference between some RCTs and observational studies is masking (blinding) of trial participants and personnel, which can be achieved by using a placebo that is indistinguishable from the active treatment. Masking prevents differential care during follow-up, accounts for nonspecific effects associated with receiving an intervention (placebo effects), may facilitate blinding of outcome assessors, and may improve adherence.

Widespread use of masking and of intention-to-treat analyses became established by regulatory requirements, which privileged intention-to-treat analyses of double-blind placebo-controlled RCTs to assess the efficacy of drugs before licensing. However, masking is sometimes not feasible (e.g., in surgical trials), and may not even be desirable (e.g., in pragmatic trials whose goal is estimating effects in real-world conditions). An intention-to-treat analysis is not feasible if trial participants are lost to follow-up and has disadvantages in safety and non-inferiority trials.²

Discussions about the differences between RCTs and observational studies can be complicated by the different terminology employed by trialists and epidemiologists.³ Trialists often use the taxonomy of bias typified by the Cochrane tool for assessing risk of bias in RCTs: selection bias, performance bias, detection bias, attrition bias, reporting bias, and other bias.^{4,5} Epidemiologists, on the other hand, tend to use the categories confounding, selection bias, and measurement (or information) bias.^{1,6,7}

Causal diagrams have been used extensively to represent biases in epidemiologic studies.⁸⁻¹⁴ These diagrams, represented as directed acyclic graphs, comprise variables (nodes) and arrows (directed edges). The absence of an arrow pointing from variable A to variable B indicates that variable A does not have a direct causal effect on B. A key advantage of causal diagrams is that they provide a mathematically rigorous yet intuitive tool for deducing the statistical independencies implied by the lack of causal arrows.^{1,8,9}

Here we use causal diagrams to represent the biases described in the Cochrane Risk of Bias Tool, and provide a translation to the epidemiologic terms of confounding, selection bias, and measurement bias. For simplicity, we focus on individually randomized (non-cluster-randomized), parallel group (non-crossover) trials that compare two time-fixed treatment strategies. We start by reviewing the main types of causal effect that are of interest in RCTs.

Intention-to-treat effect and per-protocol effect

The intention-to-treat effect is the effect of treatment assignment (or allocation).¹⁵ Consider an RCT in which HIV-positive individuals are assigned to either initiating a new treatment $Z=1$ or to continuing on their existing treatment $Z=0$, and are followed until death or the end of follow-up at 5 years. The outcome of interest is 5-year mortality Y (1: yes, 0: no). The intention-to-treat effect is unbiasedly estimated by an intention-to-treat analysis that compares the mean outcome between those assigned to $Z=1$ and $Z=0$. For example, the intention-to-treat effect on the causal risk difference scale is unbiasedly estimated by the difference of the risks in the groups $Z=1$ and $Z=0$, which are readily computed from the study data.

The magnitude of the intention-to-treat effect in a particular study depends on the magnitude and type of adherence to the assigned treatment strategies. To see this, consider two RCTs with identical eligibility criteria and that compare the same two strategies. In the first RCT, only half of the patients assigned to the new treatment ($Z=1$) end up actually taking it ($A=1$); the other half do not take it ($A=0$). In the second RCT, all patients assigned to treatment take it (that is, patients with $Z=1$ also have $A=1$). In both studies all patients assigned to $Z=0$ cannot take the new treatment because it is not available outside the study (that is, patients with $Z=0$ also have $A=0$). Even if the effect of the new treatment is identical in both studies, the intention-to-treat effect will generally differ between the two studies. For example, the intention-to-treat effect will be closer to the null in the first RCT than in the second one if the effect of treatment goes in the same direction (beneficial or harmful) for all patients, and more beneficial in the first RCT than in the second one if adherers tend to be those for whom treatment has a beneficial effect and non-adherers tend to be those for whom treatment has a harmful effect.¹ If the above RCTs were head-to-head trials that assigned participants to two active treatments, then the intention-to-treat effect in the first RCT might also be either closer or further from the null than that in the second RCT.²

An alternative to the intention-to-treat effect that is not affected by the study-specific adherence to treatment is the per-protocol effect, that is, the causal effect that would have been observed if all patients had adhered to the protocol of the RCT. Unfortunately, valid estimation of the per-protocol effect in the presence of imperfect adherence generally requires untestable assumptions.¹⁶

Two common approaches to estimate the per-protocol effect are (i) comparing the outcomes of those who took treatment $A=1$ and treatment $A=0$ (regardless of the treatment they were assigned to), e.g., $\Pr[Y=1|A=1] - \Pr[Y=1|A=0]$, and (ii) comparing the outcomes of those who took treatment $A=1$ among those assigned to $Z=1$ and treatment $A=0$ among those assigned to $Z=0$, e.g., $\Pr[Y=1|A=1, Z=1] - \Pr[Y=1|A=0, Z=0]$. Approach (i) is often referred to as an “as treated” analysis and approach (ii) as a “per protocol” analysis.^{2,8} Neither approach is generally valid to estimate the per-protocol effect, as we discuss below. (G-estimation and instrumental variable methods can sometimes be used to estimate some form of per-protocol effects even in the presence of unmeasured confounders in Figures 1c and 1d.^{16,17})

Although “as treated” and “per-protocol” analyses are potentially biased, the per-protocol effect may be of greater interest to patients and their clinicians than the intention-to-treat effect. We now discuss how the potential for bias in effects estimated from RCTs depends on whether the goal is to estimate the per-protocol or the intention-to-treat effect.

Cochrane bias domains and causal diagrams

The Cochrane Risk of Bias Tool for randomized trials covers six domains of bias.^{4,5} In the next sections, we use causal diagrams to show the structure of most of these biases, and discuss their correspondence to the epidemiologic terms of confounding, selection bias, and measurement bias. Because all these biases can occur under the null, we draw the causal diagrams under the causal null hypothesis, unless otherwise specified. (Any causal structure

that results in bias under the null will also cause bias under the alternative that treatment has an effect on the outcome, but the converse is not true.) For each bias, we explain whether it affects the intention-to-treat effect or the per-protocol effect. Our definition of bias is the same as in Chapter 10 of reference 1 under either a randomization model or a correct population model.^{18,19}

Selection bias

In its Risk of Bias Tool, Cochrane defines selection bias as the result of “systematic differences between baseline characteristics of the groups that are compared.”⁴ The presence of “systematic differences between baseline characteristics” means that the distribution of prognostic factors varies between the groups being compared. The bias may affect the estimate of the intention-to-treat effect and/or the estimate of the per-protocol effect, depending on the definition of “groups that are compared.”

Let us first consider the case in which the “groups that are compared” are the randomized groups $Z=1$ and $Z=0$. There are at least three reasons why differences in the distribution of risk factors may arise.

(i) The assignment of patients to a group is influenced by knowledge of which treatment they will receive

This bias can occur if the assignment that was not properly randomized or the randomized assignment was not sufficiently concealed, and so the person enrolling participants was aware of allocation sequence and influenced which patients were assigned to each group based on their prognostic factors. This situation is depicted by the causal diagram in Figure 1a that includes the prognostic factors L (e.g., CD4 count, viral load) as common causes of the outcome Y and the assignment Z . The arrow from L to Z may be due to the improperly randomized or insufficiently concealed allocation sequence. There are other causal diagrams that represent common causes of Z and Y (see, for example, Chapter 7 of Reference 1); we chose the simplest.

Epidemiologists refer to biases that arise from the presence of common causes as confounding. The existence of common causes L_Z of assignment Z and outcome Y introduces confounding bias for the intention-to-treat effect in an intention-to-treat analysis that compares individuals in groups $Z=1$ and $Z=0$, and for the per-protocol effect in a per-protocol analysis that compares individuals in groups $Z=1$ and $Z=0$ with $A=Z$. In both cases, the Cochrane Risk of Bias Tool refers to this bias as selection bias; see Table.

Appropriate randomization, generation, and concealment of the allocation sequence, or adjustment for the prognostic factors L removes the $L \rightarrow Z$ arrow and therefore the confounding bias.

Even under perfect randomization procedures, random imbalances in prognostic factors may bias intention-to-treat effect estimates. This so-called chance confounding²⁰ (sometimes referred to as allocation bias¹⁹ or accidental bias^{18,21,22}) is quantitatively addressed by frequentist confidence intervals and is mitigated by adjusting for measured prognostic

factors that are imbalanced.^{1,10} Unlike the structural confounding depicted in Figure 1a, chance confounding is expected to become smaller as sample size increases.

(ii) The decision to recruit a patient is influenced by knowledge of which treatment the patient will receive.^{19,21}

This bias can occur if an investigator is aware of the random sequence and decides to enroll patients with certain prognostic factors only if they are known to be assigned to a particular treatment strategy. The Cochrane Risk of Bias Tool describes this problem: “Knowledge of the next assignment [...] can cause selective enrolment of participants on the basis of prognostic factors. Participants who would have been assigned to an intervention deemed to be ‘inappropriate’ may be rejected.”⁴

The causal diagram in Figure 1b represents this scenario. The node S is the selection into the trial (1: yes, 0: no), which depends on the values of assignment Z and prognostic factors L . The box around S indicates that the analysis is restricted to those with $S=1$. This bias arises from the selection of a subset of the potential study population into the analysis and, because S is a common effect of assignment and prognostic factors, both intention-to-treat and per-protocol analyses may be biased even if both effects are truly null. Epidemiologists¹⁰ and Cochrane refer to this bias as selection bias (Table 1).

The elimination of this selection bias requires removing the $Z \rightarrow S$ arrow through appropriate concealment of the allocation sequence, or adjustment for the prognostic factors L .

(iii) The decision to adhere to the assigned treatment is influenced by prognostic factors

This may result in an imbalance between the groups $A=1$ and $A=0$, but not between the groups $Z=1$ and $Z=0$. Therefore, this imbalance will not bias the intention-to-treat estimate, but will generally bias the per-protocol estimate of a naïve per-protocol analysis. This third case is not addressed by the Cochrane Risk-of-Bias Tool.

The causal diagram in Figure 1c represents this scenario. The node U_A stands for common causes (e.g., symptoms resulting from severe immunosuppression) of adherence to treatment A and outcome Y .² The node S is the variable selection into the per-protocol population (1: yes, 0: no), which depends on the values of Z and A ($S=1$ if $Z=A$, $S=0$ if $Z \neq A$), and the analysis is restricted to those with $S=1$.

Epidemiologists may refer to this bias as selection bias because it arises from the selection of a subset (the per-protocol population) of the study population into an analysis that compares $Z=1$ vs. $Z=0$. However, note that this selection bias for the per-protocol effect only arises when there is confounding of the effect of A due to common causes U_A of A and Y .

Regardless of whether we refer to this bias as confounding or selection bias, reducing it requires either a masked design or a non-naïve, more realistic per-protocol analysis that adjusts for the variables U_A or their proxies. Because a per-protocol analysis compares groups not entirely defined by randomization, the analysis is subject to the biases usually associated with observational studies.

Finally, let us consider the case in which the “groups that are compared” are the non-randomized groups $A=1$ and $A=0$, i.e., an as-treated analysis. Because as-treated analyses are effectively observational analyses, their estimates of the per-protocol effect will be confounded in the presence of common causes of treatment A and the outcome Y . The structure of the bias is shown by the causal diagram in Figure 1d. In subsequent diagrams we will omit the common causes of A and Y to avoid clutter and to focus the attention on the other sources of bias.

Performance bias

The Cochrane Risk of Bias Tool defines performance bias as the result of “systematic differences between groups in the care that is provided, or in exposure to factors other than the interventions of interest.”³ These differences may occur due to knowledge of the assigned treatment Z by study participants and thus will be less likely in masked trials.

Again let us first consider the case in which the “groups that are compared” are the randomized groups $Z=1$ and $Z=0$. Consider the causal diagram in Figure 2a, where O represents medical interventions that are forbidden by the study protocol (e.g., an intensive monitoring and treatment of cardiovascular risk factors) and U_O represents unmeasured common causes of O and Y (e.g., risk factors for cardiovascular disease). The arrow from Z to O indicates that awareness of the assigned treatment might lead to changes in the behavior of study participants or their doctors, which in turn may affect the outcome, hence an arrow from O to Y . The interventions O are a result of assignment Z itself, and therefore just mediators of the effect of Z .

Because the intention-to-treat effect is the effect of assignment and part of the effect of assignment is mediated through O , then O cannot be viewed as a source of bias. The intention-to-treat effect naturally incorporates the effects of deviations from protocol, including the interventions O . That is, in an intention-to-treat analysis whose goal is to estimate the intention-to-treat effect, performance bias cannot occur. In epidemiologic terms, there is no confounding or selection bias.

However, the Cochrane literature appears to suggest that performance bias can occur even in intention-to-treat analyses. To explore this issue, consider two types of departures from intended interventions

(i) Departures from intended interventions that might happen in real life

When trial participants receive interventions O that are prohibited by the protocol but that they would have also received outside of the trial, we believe that most people would agree with the conclusion that no performance bias exists in intention-to-treat analyses.

(ii) Departures from intended interventions that arise only because of the randomized trial context

When trial participants receive interventions O that are prohibited by the protocol and that they would have not received outside of the trial, the intention-to-treat effect estimated from the trial is relatively unhelpful for patients outside the trial. This may be a reason why

Cochrane uses the performance bias label for randomized trials. However, the use of the word “bias” in this context needs to be carefully qualified.

In the absence of any of the other biases discussed here, an intention-to-treat analysis of a trial in which interventions O occur is an unbiased estimator of the intention-to-treat effect in that particular trial context and population. The presence of interventions O , like other trial-specific characteristics (e.g., eligibility criteria, monitoring) do not affect the estimates' internal validity but that may affect their external validity (e.g, if the estimate cannot be transported to clinical contexts outside of the trial in which the interventions O are less frequent). In this case, it might be more appropriate to say that the intention-to-treat effect from the trial is not generalizable or transportable to other settings rather than saying that it is “biased”.

Another reason why Cochrane uses the performance bias label for intention-to-treat analysis of randomized trials is that the implicit research question may not be about the pure intention-to-treat effect, but rather an intention-to-treat effect where the only deviation in protocol is non-adherence to the assigned treatment i.e., the research question is neither intention-to-treat effect nor per-protocol effect. Thus, different research questions might lead to different categorizations of bias.

Performance bias may occur when estimating the per-protocol effect via either a per-protocol or an as-treated analysis. There are at least two distinct reasons for the bias to arise.

Figure 2b depicts a setting in which deviations from protocol O are affected by the received treatment (e.g., because the use of certain therapies prompts doctors conduct tests for cardiovascular risk factors that were forbidden by the protocol). The per-protocol effect is then the direct effect of treatment in the absence of those deviations from protocol O (e.g., if no tests for cardiovascular factors had been conducted). Unfortunately, per-protocol and as-treated analyses will yield a biased direct effect (per-protocol) estimate whether they do or do not adjust for O (Table 1). Lack of adjustment will result in bias because the effect estimate will include the indirect effect as well; adjustment for O but not also for all confounders of the effect of O on the outcome (e.g., U_O),¹¹ will generally result in selection bias (in graph theoretic terms, O is a collider).

Performance bias for the per-protocol effect may also have the same structure as confounding. Figure 2c represents a setting in which O operates as a confounder of the effect of A on Y . The arrow from O to A indicates that interventions not specified in the protocol (e.g., intensive monitoring of cardiovascular risk factors) may alter treatment received during the follow-up (e.g., the presence of cardiovascular risk factors leads doctors to prescribe a different type of antiretroviral therapy). With time-varying variables A and O , the structures represented in Figures 2b and 2c can occur simultaneously.

Detection bias

The Cochrane Risk-of-Bias Tool defines detection bias as the result of “systematic differences between groups in how outcomes are determined”.⁴ This bias (also called observer, ascertainment, or assessment bias) occurs if knowledge of a patient's assigned

strategy influences outcome assessment. Figure 3a represents detection bias for the intention-to-treat effect. In this graph the true outcome Y remains unmeasured and Y^* represents the mismeasured outcome. The arrows from Z and Y to Y^* represent that outcome measurement depends on both the true outcome Y and the treatment assignment Z . An intention-to-treat estimate of the effect of Z on Y^* from Figure 3a will be biased for the intention-to-treat effect of Z on Y ; the bias is a consequence of mismeasurement of Y , and is commonly referred to as “measurement bias” or “information bias” in epidemiology (Table 1).^{1,6,12} The type of measurement error represented in Figure 3a is differential with respect to treatment assignment^{1,5} and therefore, like all other biases discussed previously, leads to bias even if Z has no effect on Y .

Detection bias may affect per-protocol effect estimates either directly if A affects Y^* , as in Figure 3b, or indirectly, as in Figure 3c, if Z is a common cause of A and Y^* . Measurement bias in Figures 3a-c can be avoided by masking of outcome assessors, because it removes the $Z \rightarrow Y^*$ arrow or the $A \rightarrow Y^*$ arrow.

Attrition bias

The Cochrane Risk of Bias Tool defines attrition bias as the result of “systematic differences between groups in withdrawals from a study”.⁴ The source of bias is differential loss-to-follow-up (e.g., drop out) or other forms of exclusions from the analysis. Figure 4a includes the censoring indicator C , which takes value 1 for individuals excluded from the analysis. The box around $C=0$ indicates that the analysis is restricted to those who were not excluded from the analysis. The arrow from Z to C indicates that withdrawal from the analysis is influenced by knowledge of the participant's group assignment, e.g., patients assigned to less potent combination antiretroviral therapy are more likely to not attend future visits if they were aware of their assigned treatment. The arrow from L to C indicates that individuals with worse prognosis ($L=1$) are more likely to be excluded than the others (with $L=0$), because the severity of their disease prevents them from attending future study visits. In graph-theoretic terms, the intention-to-treat effect estimate is biased because, even under the null, the path $Z \rightarrow C \leftarrow L \rightarrow Y$ is open when conditioning on the collider C . This bias is another example of what epidemiologists refer to as selection bias (Table 1).^{1,10}

The per-protocol effect estimate is also subject to attrition bias. The bias may arise directly if A affects censoring C (e.g., subjects receiving $A=1$ are at a greater risk of experiencing side effects, which could lead them to dropout), as in Figure 4b, or indirectly, as in Figure 4c, if Z is a common cause of A and C . In Figures 4a and 4c, masking of participants and doctors providing care can prevent attrition bias for intention-to-treat and per-protocol effect estimates by removing the $Z \rightarrow C$ arrow, and thus blocking the biasing paths $Z \rightarrow C \leftarrow L \rightarrow Y$ and $A \leftarrow Z \rightarrow C \leftarrow L \rightarrow Y$, respectively. Adjustment for L also in Figures 4a-c also adjusts for selection bias.

Reporting bias

The Cochrane Risk of Bias Tool defines reporting bias as the result of “systematic differences between reported and unreported findings”.⁴ Outcome reporting bias may occur

because so-called statistically significant effect estimates are more likely to be reported than nonsignificant effect estimates. Then the average published result will be farther from the null than the true average result, which will bias meta-analyses and systematic reviews.²³ A similar bias, “stepwise selection”, results in inflated estimated for weak effects, sometimes known as testimation (estimation after testing) bias.^{24–26} Reporting bias, which applies to both intention-to-treat and per-protocol effects, is negligible when treatment has a strong effect on the outcome or the trial size is huge.^{24,25}

Epidemiologists have long warned against the problems resulting from abuse of significance testing and selective reporting after multiple comparisons.^{27,28} Because reporting bias of individual studies is not a structural bias, it cannot be generally represented using the causal diagrams.

Discussion

We described how the terminology used to describe similar biases differs between trialists and epidemiologists, and why an explicit specification of the causal target of each randomized trial is beneficial when discussing the risk of bias. For example, making the intention-to-treat the target allows trialists to stop worrying about some forms of “performance bias”. On the other hand, making the per-protocol effect the target makes it clear that adjustment for pre- and post-randomization confounding is needed, which has implications for the design and analysis of RCTs.²⁹

We encourage trialists and epidemiologists to be more explicit about their inferential goals. In particular, an open question is whether trialists conducting intention-to-treat analyses are really interested in the intention-to-treat effect. The widespread preference for masked studies suggests that the per-protocol effect, which is not affected by differential implementation of the treatment strategies being compared, may be the ultimate target.

Causal diagrams help reduce confusion created by ambiguous terminology.³⁰ For example, the term selection bias is used with different meanings by trialists and epidemiologists. Drawing the corresponding causal diagram helps resolve these confusions. The structural approach to bias using causal diagrams also shows that some biases that are described using different terms in the RCT literature have the same structure. For example, Figures 1b and 4a are essentially the same apart from the time and reason for selection.

Our simplistic graphical presentations of the Cochrane selection, performance, detection and attrition biases cannot possibly cover all possibilities. Specifically, in trials with time-varying treatments and attrition, Robins's g-methods (g-formula, inverse-probability weighting, g-estimation) are generally needed to properly adjust for time-varying confounding and selection bias when estimating intention-to-treat and per-protocol effects.¹⁶

Acknowledgments

Funding: This work was funded by NIH grant R01 AI102634, and a Methods Innovation Fund grant from the Cochrane Collaboration. Jonathan Sterne is funded by National Institute for Health Research Senior Investigator award NF-SI-0611-10168.

References

1. Hernán, MA.; Robins, JM. Causal Inference. London: Chapman & Hall/CRC; 2016.
2. Hernán MA, Hernández-Díaz S. Beyond the intention to treat in comparative effectiveness research. *Clinical Trials*. 2012; 9:48–55. [PubMed: 21948059]
3. Heitjan DF. Ignorability and bias in clinical trials. *Stat Med*. 1999; 18:2421–34. [PubMed: 10474150]
4. Higgins, JPT.; Green, S., editors. *Cochrane Handbook for Systematic Reviews of Interventions*. Chichester: John Wiley & Sons; 2008.
5. Higgins JPT, Altman DG, Gøtzsche PC, Jüni P, Moher D, Oxman AD, et al. The Cochrane Collaboration's tool for assessing risk of bias in randomised trials. *BMJ*. 2011; 343:d5928. [PubMed: 22008217]
6. Rothman, KJ.; Greenland, S.; Lash, TL. Validity in epidemiologic studies. In: Rothman, KJ.; Greenland, S.; Lash, T., editors. *Modern Epidemiology*. 3rd. Philadelphia, PA: Lippincott Williams & Wilkins; 2008. p. 128-147.
7. Miettinen, OS. *Theoretical epidemiology*. New York: Wiley; 1985.
8. Pearl, J. *Causality*. 2nd. New York: Cambridge University Press; 2009.
9. Greenland S, Pearl J, Robins JM. Causal diagrams for epidemiologic research. *Epidemiology*. 1999; 10:37–48. [PubMed: 9888278]
10. Hernán MA, Hernandez-Diaz S, Robins JM. A structural approach to selection bias. *Epidemiology*. 2004; 15:615–625. [PubMed: 15308962]
11. Cole SR, Hernán MA. Fallibility in estimating direct effects. *Int J Epidemiol*. 2002; 31:163–5. [PubMed: 11914314]
12. Hernán MA, Cole SR. Invited Commentary: Causal diagrams and measurement bias. *Am J Epidemiol*. 2009; 170:959–62. [PubMed: 19755635]
13. Hernan MA, Hernandez-Diaz S, Werler MM, Mitchell AA. Causal knowledge as a prerequisite for confounding evaluation: An application to birth defects epidemiology. *Am J Epidemiol*. 2002; 155:176–184. [PubMed: 11790682]
14. Mansournia MA, Hernán MA, Greenland S. Matched designs and causal diagrams. *Int J Epidemiol*. 2013; 42:860–869. [PubMed: 23918854]
15. Food and Drug Administration. International Conference on Harmonisation; Guidance on Statistical Principles for Clinical Trials. Federal Register. 1998; 63:49583–49598. [PubMed: 10185190]
16. Toh S, Hernán MA. Causal inference from longitudinal studies with baseline randomization. *Int J Biostat*. 2008; 4 Article 22.
17. Hernán MA, Robins JM. Instruments for causal inference—An epidemiologist's dream? *Epidemiology*. 2006; 17:360–372. [PubMed: 16755261]
18. Rosenberger, WF.; Lachin, JM. *Randomization in Clinical Trials: Theory and Practice*. New York: Wiley; 2002.
19. Matthews, JNS. *An Introduction to Randomised Controlled Clinical Trials*. 2nd. Boca Raton: Chapman & Hall/CRC; 2006.
20. Greenland S, Mansournia MA. Limitations of individual causal models, causal graphs, and ignorability assumptions, as illustrated by random confounding and design unfaithfulness. *Eur J Epidemiol*. 2015; 30:1101–10. [PubMed: 25687168]
21. Friedman, LM.; Furberg, CD.; DeMets, DL. *Fundamentals of Clinical Trials*. 4th. New York: Springer; 2010.
22. Chow, SC.; Liu, JP. *Design and Analysis of Clinical Trials*. 2nd. Hoboken, NJ: Wiley-Interscience; 2004.
23. Greenland, S.; O'Rourke, K. Meta-analysis. In: Rothman, KJ.; Greenland, S.; Lash, T., editors. *Modern Epidemiology*. 3rd. Philadelphia, PA: Lippincott Williams & Wilkins; 2008. p. 652-682.
24. Steyerberg, EW. *Clinical prediction models: a practical approach to development, validation, and updating*. New York, NY: Springer; 2008.

25. Steyerberg EW, Eijkemans MJ, Habbema JD. Stepwise selection in small data sets: a simulation study of bias in logistic regression analysis. *J Clin Epidemiol*. 1999; 52:935–942. [PubMed: 10513756]
26. Maldonado G, Greenland S. Simulation study of confounder-selection strategies. *Am J Epidemiol*. 1993; 138:923–936. [PubMed: 8256780]
27. Greenland, S.; Rothman, KJ. Fundamentals of epidemiologic data analysis. In: Rothman, KJ.; Greenland, S.; Lash, T., editors. *Modern Epidemiology*. 3rd. Philadelphia, PA: Lippincott Williams & Wilkins; 2008. p. 213-237.
28. Ioannidis JPA. Why most discovered true associations are inflated. *Epidemiology*. 2008; 1:640–648.
29. Hernán MA, Hernández-Díaz S, Robins JM. Randomized trials analyzed as observational studies. *Ann Intern Med*. 2013; 159:560–2. [PubMed: 24018844]
30. Shrier I. Structural approach to bias in meta-analyses. *Res Synth Methods*. 2011; 2:223–37. [PubMed: 26061887]

Figure 1a. Intention-to-treat effect: Cochrane selection bias/Epidemiologic confounding in an intention-to-treat analysis

Figure 1b. Intention-to-treat effect: Cochrane selection bias/Epidemiologic selection bias in an intention-to-treat analysis

Figure 1c. Per-protocol effect: Cochrane selection bias/ Epidemiologic selection bias in a per-protocol analysis

Figure 1d. Per-protocol effect: Cochrane selection bias/ Epidemiologic confounding in an as-treated analysis

Figure 2a. Intention-to-treat effect: Cochrane RCT performance bias is not a bias in intention-to-treat analyses

Figure 2b. Per-protocol effect: Cochrane performance bias/ Biased direct effect in a per-protocol or as-treated analysis

Figure 2c. Per-protocol effect: Cochrane performance bias / Epidemiologic confounding a in per-protocol or as-treated analysis

Figure 3a. Intention-to-treat effect: Cochrane detection bias/ Measurement bias due to outcome misclassification in an intention-to-treat analysis

Figure 3b. Per-protocol effect: Cochrane detection bias/ Measurement bias due to outcome misclassification in a per-protocol or as-treated analysis

Figure 3c. Per-protocol effect: Cochrane detection bias/ Measurement bias due to outcome misclassification in a per-protocol or as-treated analysis

Figure 4a. Intention-to-treat effect: Cochrane attrition bias/ Selection bias due to differential loss-to-follow-up in an intention-to-treat analysis

Figure 4b. Per-protocol effect: Cochrane attrition bias/ Selection bias due to differential loss to follow-up in a per-protocol or as-treated analysis

Figure 4c. Per-protocol effect: Cochrane attrition bias/ Selection bias due to differential loss to follow-up in a per-protocol or as-treated analysis

Table
Translation of Cochrane bias in randomized trials domains into common epidemiologic terms

Cochrane bias domain	Epidemiologic term	Bias in intention-to-treat effect	Bias in per-protocol effect
Selection bias	Confounding or selection bias	Yes	Yes
Performance bias	Biased direct effect or confounding	No	Yes
Detection bias	Measurement bias	Yes	Yes
Attrition bias	Selection bias	Yes	Yes
Reporting bias	Non-structural bias that cannot be represented in our causal diagrams	Yes	Yes