Published in final edited form as:

Epidemiology. 2018 July ; 29(4): e24–e25. doi:10.1097/EDE.0000000000000823.

On well-defined hypothetical interventions in the potential outcomes framework

Tyler J. VanderWeele

¹Departments of Epidemiology and Biostatistics, Harvard T.H. Chan School of Public Health, Harvard University, Boston, Massachusetts

A recent series of articles, commentaries, letters, and responses across numerous journals¹⁻¹² has considered the role and limits of the potential outcomes framework within the broader undertaking of causal inference. In one response in that exchange Broadbent et al.⁷ noted that the potential outcomes framework often makes reference to the notion of a “hypothetical intervention” but typically without formal explication of what this is. Often when the potential outcomes framework is described, the potential outcomes themselves, Y_1 and Y_0 , are understood as the outcomes that would occur under a “hypothetical intervention” to set the exposure or treatment to levels 1 and 0 respectively. The term “hypothetical intervention” is generally taken as a primitive and left undefined.

I would like to put forward a proposal here as to how a “hypothetical intervention”, within the potential outcomes framework, is to be understood. I am not suggesting that there is anything new in this proposal, and indeed I think what is here proposed is what is generally implicitly understood by many writers on the potential outcomes framework when using the term. The hope of this letter is thus only to make explicit what I think is often understood implicitly, and thus, to open the proposal to additional comment, refinement and critique. I will be adapting here some discussion in Chapter 16 of my book,¹³ *Explanation in Causal Inference*, that relates causation to entailment by the laws of nature. Even in my book, while I certainly had a notion of what constituted a “hypothetical intervention” in mind, and discussed what I took to be the relevant concepts in reference to statements of the form “X had an effect on Y”, I did not explicitly tie this to the notion of a “hypothetical intervention,” which again often seems to be taken as a primitive, left undefined, within the potential outcomes framework.

I would propose that a hypothetical intervention to set X to x is simply the specification, possibly contrary to fact, of the event or state $X=x$, such that we say that a hypothetical intervention is well-defined with respect to outcome Y, exposure X, setting $X=x$, and population P, if for each individual i in population P, there is a unique value $Y_x(i)$ (or distribution of values $Y_x(i)$ in the context of stochastic counterfactuals) such that the event or state $X=x$, along with the state of the universe and the laws of nature, jointly entail $Y=Y_x(i)$. Otherwise, the hypothetical intervention is said to be ill-defined or ambiguous. A hypothetical intervention on X is thus essentially just the specification of a potentially

counterfactual state of exposure X but one such that the state or event $X=x$ is itself sufficiently well-specified so as to fix the value or distribution Y_x for each individual in the population under consideration. We might also refer to a “hypothetical intervention” as simply a “hypothetical setting” or, when the setting is contrary to fact, a “counterfactual state.”

As discussed in the eAppendix, the ambiguity of hypothetical intervention potentially arises from failure to sufficiently specify the state of the universe that is in view at least as it pertains to the outcome under consideration. The extent to which a hypothetical intervention is well defined or ambiguous is of course a relative concept. Most specifications of an event or state will be ambiguous to a certain extent, but again for the hypothetical intervention to be well defined the specification of the state only needs to be sufficiently fixed so as to entail a unique outcome or distribution of outcomes Y . When hypothetical interventions are well defined, then so are the potential outcomes, and thus so are also the quantitative causal effect estimands.⁴ Discussion is also given in the eAppendix as to challenges in conceptualizing the laws of nature themselves, how the proposal here relates to Pearl’s “do” operator, how it relates to multiple versions of treatment, and how it helps clarify that manipulability is helpful, but not necessary, in specifying well-defined hypothetical interventions.

I hope that this discussion here and in the eAppendix goes some way in addressing Broadbent et al.’s⁷ inquiry concerning the conceptualization of “hypothetical interventions” in the potential outcomes literature, and that it will also allow for further discussion and critique.

eAppendix for “On well-defined hypothetical interventions in the potential outcomes framework” by Tyler J. VanderWeele

The eAppendix here explores further the implications of the proposal in the text concerning hypothetical interventions being well defined. See other work elsewhere for related discussion.^{1–19} As noted in the text, I would propose that a hypothetical intervention to set X to x is simply the specification, possibly contrary to fact, of the event or state $X=x$, such that we say that a hypothetical intervention is well-defined with respect to outcome Y , exposure X , setting $X=x$, and population P , if for each individual i in population P , there is a unique value $Y_x(i)$ (or distribution of values $Y_x(i)$ in the context of stochastic counterfactuals^{20,21}) such that the event or state $X=x$, along with the state of the universe and the laws of nature, jointly entail $Y=Y_x(i)$. Otherwise, the hypothetical intervention is said to be ill-defined or ambiguous. A hypothetical intervention on X is thus essentially just the specification of a potentially counterfactual state of exposure X but one such that the state or event $X=x$ is itself sufficiently well-specified so as to fix the value or distribution Y_x for each individual in the population under consideration. We might also refer to a “hypothetical intervention” as simply a “hypothetical setting” or, when the setting is contrary to fact, a “counterfactual state.”

Consider a single individual i . For this individual, a hypothetical intervention would then not be well-defined (i.e. we would say there is ambiguity in the hypothetical intervention) if there are multiple states of the universe consistent with the specified event or state $X=x$, and

if some of these states of the universe, along with the event $X=x$ entail different values of (or distributions) of $Y(i)$. For a population, the hypothetical intervention will be ambiguous if it is ambiguous for any individual i in the population. The hypothetical intervention is well-defined for the population if and only if it is well-defined for each individual.

Thus, although there may be multiple states of the universe consistent with the state or event $X=x$, it is possible that all of these would entail the same value or distribution of Y_x , and then the hypothetical intervention is still unambiguous. In other cases, there may be multiple states of the universe consistent with $X=x$, and it may be the case that the laws of nature do not entail a unique value or distribution Y_x , but if it is further clarified which state or states of the universe are in view, at least to the extent that the laws of nature and the state or states of the universe thus specified entail a unique value or distribution Y_x , then, once again, the hypothetical intervention would be well-defined. In this case we would say that the hypothetical intervention has been specified further, and is no longer ambiguous^{4,13,17,22}. We can thus, to a certain extent, clarify the ambiguity of a hypothetical intervention by further specifying the state or states of the universe in view, at least to the extent necessary for there to be a unique value or distribution of Y_x entailed by $X=x$, the state of the rest of the universe, and the laws of nature.

In summary, we might thus conceptualize a hypothetical intervention in terms of entailment, of an outcome or distribution, by the state of the universe and the laws of nature. The ambiguity of hypothetical intervention again then potentially arises from failure to sufficiently specify the state of the universe that is in view at least as it pertains to the outcome under consideration. When hypothetical interventions are well-defined, then so are the potential outcomes, and thus so are also the quantitative causal effect estimands.⁴

It is my belief that the description above roughly coincides with how the term “hypothetical intervention” is often being employed in the potential outcomes literature. The above description merely tries to formalize current usage. It does not resolve all philosophical difficulties as appeal is made to entailment by the laws of nature and the status of what these are is of course also disputed. Further discussion of these points is given elsewhere^{14,17,18,23–26}. Very roughly, the laws of nature might be understood as a minimal set of statements that describe the regularities among physical phenomena and that would hold even in contrary-to-fact states of the universe; science seeks to catalogue these laws of nature, by induction from hypothesis and repeated experimentation in attempt to obtain a minimal set of regularities with maximum explanatory power. Making such a description of the laws of nature more precise is the subject of a great deal of philosophical and scientific literature. For the purposes of the proposal above, it is taken as a concept that is more fundamental than that of a “well-defined hypothetical intervention” and from which the latter is derivative.

This description of a hypothetical intervention also, I believe, roughly corresponds to Pearl’s “do” operator, or “surgery on equations” in which there is no articulation of how the exposure was set^{27,28}, but again, for this to make sense, the state $X=x$ must be sufficiently well-defined so as to entail, along with the rest of the statement of the universe (captured in Pearl’s model by the other nodes on the diagram) a unique distribution⁴.

Importantly here, the “hypothetical intervention” can be a state or event that need not correspond to an actual intervention⁹, and thus moreover, it need not correspond to a realistic practical intervention, a point often made by those employing the framework, but disputed by a minority²⁹; cf. 7. For example, with the definitions above, we could speak of a well-defined hypothetical intervention on a genetic variant even if we cannot at present alter it. We can do this because it is relatively straightforward to think of just that one variant being other than it was and the rest of the universe being the same.

In contrast, if we were to consider the temperature being 30 degrees versus 40 degrees in a specific town, there may have to be numerous things that would have to be different for the temperature to be different, such as the wind speed and air pressure in that town, along with the temperature in spatially contiguous towns, etc. There may be multiple states of the universe that are consistent with the temperature in a specific town being 30 degrees rather than 40 degrees, and these different states of the universe may have very different implications for the outcome Y. Here we would have considerable ambiguity in the “hypothetical intervention” or counterfactual state being specified.

Sometimes it seems we are more comfortable with hypothetical interventions grounded in human action perhaps because it seems to be the case that human actions, at least on the surface of things, seem sufficiently free that we have an easier time imagining only one specific action being different, and nothing else. With regard to what else would have had to be different if a human action had been different, there seems to be less ambiguity than there is if some other aspect of the state of the universe had been different. It is easier to imagine the rest of the universe being just as it is if a patient took pill A rather than pill B than it is trying to imagine what else in the universe would have had to be different if the temperature had been 30 degrees rather than 40 degrees.¹⁷ A hypothetical intervention grounded in human action can thus perhaps often help make a hypothetical intervention less ambiguous, but it is not necessary, as the example of the genetic variant indicates. Moreover, it is also not sufficient insofar as there can be ambiguity in the action itself (e.g. whether the pill A is taken in the morning or in the evening) and if that ambiguity is relevant for the outcome for certain individuals, then the hypothetical intervention, even though it consists of a human action, is still ambiguous. Manipulability versus well-defined hypothetical intervention (and thus also quantitative causal effects estimands) are thus related, but non-overlapping, domains.

We might speak of a hypothetical intervention grounded in human action to change an exposure as a “manipulable exposure or treatment”^{7,29}, or a practicable intervention⁹. This is a subclass of the hypothetical interventions described above, but again one that does not itself necessarily guarantee that the resulting hypothetical intervention is well-defined. While there is likely less that needs to be considered about what else in the state of the universe must be different for the human action to be different than with the specification of many counterfactual states not grounded in human action, this again, does not guarantee that all ambiguity is eliminated. That a hypothetical intervention is practicable, or concerns a manipulable exposure, is thus more closely related to whether it is useful from a policy or interventionist perspective, than to whether the causal inquiry is itself well-defined.

The extent to which a hypothetical intervention is well-defined or ambiguous is of course a relative concept. Most specifications of an event or state will be ambiguous to a certain extent, but again for the hypothetical intervention to be well-defined the specification of the state only needs to be sufficiently fixed so as to entail a unique outcome or distribution of outcomes Y . A position is often articulated in the potential outcomes literature that all counterfactuals are vague to a certain extent, but that they are made more precise by specifying further the relevant contrary-to-fact scenario or by specifying further what is meant by the counterfactual^{4,13,17,18,30}. This also is consistent with the definitions suggested above.

The level of ambiguity that an investigator is willing to tolerate will likely vary by context and outcome, and, as clear in the definition above, whether the hypothetical intervention is well defined is relative to the outcome under consideration. More coarsened versions of the outcome (e.g. systolic blood pressure above or below 130, rather than exact systolic blood pressure; or systolic blood pressure with no decimal places, versus systolic blood pressure rounded to the nearest multiple of ten) may also reduce the ambiguity of a hypothetical intervention with respect to the outcome thus defined. Often investigators will live with some degree of ambiguity, but understanding the extent of the ambiguity and having a more formal framework within which to reason about the extent of ambiguity with regard to a hypothetical intervention may help clarify both the extent of the ambiguity and potential ways to reduce that ambiguity.

One final comment may merit some attention. Recent theoretical work on causal inference under multiple versions of treatment^{14,15} may also provide an additional approach to interpret estimates of contrasts of potential outcomes when the relevant hypothetical interventions are still ambiguous. Discussion of these issues also takes place under violations of the so-called Stable Unit Treatment Value Assumption (SUTVA)^{31,32} and the “consistency assumption”^{19,27,30,33} that for each i , $Y_x(i)=Y(i)$ whenever $X=x$. When there are multiple versions of treatment this assumption is violated, because the potential outcomes Y_x are not well-defined since $X=x$ does not specify the version of treatment and different versions of treatment may lead to different outcomes. The recent theory for causal inference under multiple versions of treatment concerns settings in which a treatment or exposure of interest is such that the exposure itself in fact more precisely consists of a composite of different more specific versions of that exposure, so that potential outcomes indexed by the actual exposure variable used in the analysis correspond to an ill-defined hypothetical intervention but settings of the underlying versions of exposure correspond to well-defined hypothetical interventions. Under various assumptions, the estimate for the effect of the composite exposure (corresponding to an ill-defined hypothetical intervention) can be conceived of as a weighted average of the effects of the different, more precise, “versions of the exposure”, weighted by the probability of each version naturally arising within the subpopulation that actually had a particular exposure level. This interpretation under multiple versions of treatment is somewhat more difficult when the underlying distribution of the “versions of exposure” are unknown, and then the no unmeasured confounding assumptions required for this interpretation are also difficult to assess when the different underlying versions of treatment or exposure are unknown. Nevertheless, this approach offers yet another alternative way to conceive of potential outcomes when the

settings of the exposure variable used in an analysis correspond to ill-defined hypothetical interventions.

It is hoped that this discussion here goes some way in addressing Broadbent et al.'s⁷ inquiry concerning the conceptualization of “hypothetical interventions” in the potential outcomes literature, and that it will also allow for further discussion and critique.

References

1. Vandembroucke JP, Broadbent A, Pearce N. Causality and causal inference in epidemiology: the need for a pluralistic approach. *International Journal of Epidemiology*, 45(6): 1776–1786.
2. Krieger N, Davey Smith G. The tale wagged by the DAG: broadening the scope of causal inference and explanation for epidemiology. *International Journal of Epidemiology*, 45(6): 1787–1808.
3. VanderWeele TJ, Hernán MA, Tchetgen Tchetgen, E.J., and Robins JM Re: Causality and causal inference in epidemiology: the need for a pluralistic approach. *International Journal of Epidemiology*, 45(6): 2199–2200.
4. VanderWeele TJ, On Causes, Causal Inference, and Potential Outcomes. *International Journal of Epidemiology*, 45(6): 1809–1816
5. Robins JM, Weissman MB. Counterfactual causation and streetlamps: what is to be done? *International Journal of Epidemiology*, 45(6): 1830–1835.
6. Daniel RM, De Stavola BL, Vansteelandt S. The formal approach to quantitative causal inference in epidemiology: misguided or misrepresented? *International Journal of Epidemiology*, 45(6): 1817–1829.
7. Broadbent A, Vandembroucke JP, Pearce N. Formalism or pluralism? A reply to commentaries on ‘Causality and causal inference in epidemiology.’ *International Journal of Epidemiology*, 45(6): 1841–1851.
8. Hernán MA, Does water kill? A call for less casual causal inferences. *Annals of Epidemiology*, 2016;26(10):674–680. [PubMed: 27641316]
9. Schwartz S, Campbell UB, Gatto BM. Causal identification: reclaiming an abandoned charge of epidemiology. *Annals of Epidemiology*, 2016;669–673. [PubMed: 27237595]
10. Greenland S. For and Against Methodologies: Some Perspectives on Recent Causal and Statistical Inference Debates. *Eur J Epidemiol*, 2017 1;32(1):3–20. [PubMed: 28220361]
11. Schwartz S Commentary: On the Application of Potential Outcomes-based Methods to Questions in Social Psychiatry and Psychiatric Epidemiology. *Social Psychiatry and Psychiatric Epidemiology*, 2017;52:139–142. [PubMed: 28180929]
12. VanderWeele TJ Causal effects of religious service attendance? *Social Psychiatry and Psychiatric Epidemiology*, in press.
13. VanderWeele TJ. *Explanation in Causal Inference: Methods for Mediation and Interaction*. New York: Oxford University Press, 2015.
14. VanderWeele TJ and Hernán MA Causal inference under multiple versions of treatment. *Journal of Causal Inference*, 2013;1:1–20. [PubMed: 25379365]
15. Hernán MA, VanderWeele TJ Compound treatments and transportability of causal inference. *Epidemiology*, 2011;22:368–377. [PubMed: 21399502]
16. VanderWeele TJ and Hernán MA Causal effects and natural laws: towards a conceptualization of causal counterfactuals for non-manipulable exposures with application to the effects of race and sex In: Berzuini C, Dawid P, and Bernardinelli L (Eds.): *Causality: Statistical Perspectives and Applications*. Wiley Chapter 9, p. 101–113, 2012.
17. Robins JM, Greenland S Comment on “Causal inference without counterfactuals.” *Journal of the American Statistical Association*, 2000;95:477–482.
18. Lewis D *Counterfactuals*. Cambridge, MA: Harvard University Press, 1973.
19. VanderWeele TJ Concerning the consistency assumption in causal inference. *Epidemiology*, 2009;20:880–883. [PubMed: 19829187]

20. Robins JM and Greenland S The probability of causation under a stochastic model for individual risk. *Biometrics* 1989;45, 1125–38. [PubMed: 2611320]
21. VanderWeele TJ and Robins JM Stochastic counterfactuals and stochastic sufficient causes. *Statistica Sinica*, 2012; 22:379–392. [PubMed: 25473251]
22. Hernán MA, Taubman SL Does obesity shorten life? The importance of well defined interventions to answer causal questions. *International Journal of Obesity* 2008; 32:S8–S14.
23. Mill JS *A System of Logic: Ratiocinative and Inductive*. London: Longmans, Green, 1911.
24. Ramsey FP *Universal of law and of fact* In: Mellor DH, editor. *Foundations: Essays*. New York: Columbia University Press, 1928.
25. Psillos S *Causation and Explanation*. Montreal: McGill-Queens University Press, 2000.
26. Psillos S Regularity theories In: Beebe H, Hitchcock C, and Menzies P, editors. *Oxford Handbook of Causation*. Oxford: Oxford University Press, 2009, pp. 619–631.
27. Pearl J *Causality: Models, Reasoning, and Inference*, 2nd ed Cambridge: Cambridge University Press, 2009.
28. Pearl J The Foundations of Causal Inference. *Sociological Methodology* 2010;40(1):75–149.
29. Holland PW Statistics and Causal Inference. *Journal of the American Statistical Association* 1986;81:945–960.
30. Hernán MA, Robins JM *Causal Inference*. Chapman Hall, forthcoming.
31. Rubin DB Comment on: “Randomization analysis of experimental data in the Fisher randomization test” by Basu. D *Journal of the American Statistical Association* 1980; 75:591–593.
32. Rubin DB Which ifs have causal answers? Comment on: “Statistics and causal inference” by Holland. P *Journal of the American Statistical Association* 1986; 81:961–962.
33. Cole SR, Fangakis CE The consistency statement in causal inference: a definition or an assumption? *Epidemiology*, 2009; 20:3–5. [PubMed: 19234395]