

Macartan Humphreys (Ph.D., Harvard, 2003) works on the political economy of development and formal political theory. Ongoing research focuses on civil wars, post-conflict development, ethnic politics, natural resource management, political authority and leadership, and democratic development with a current focus on the use of field experiments to study democratic decision-making in post-conflict and developing areas. He has conducted field research in Chad, Ghana, Haiti, Indonesia, Liberia, Mali, Sao Tome and Principe, Sierra Leone, Senegal, Uganda, and elsewhere. Recent work has appeared in the American Political Science Review, World Politics, the Economic Journal, and elsewhere. He has authored or coauthored books on ethnic politics, natural resource management, and game theory and politics. A former Trudeau fellow and scholar of the Harvard Academy, he is a Professor of Political since at Columbia University. Humphreys was a founding member of EGAP and served as Executive Director from 2012 to 2015.

FOLLOW US



LISTSERVE

Sign up to be on our mailing list



INFORMATION

[About us \(/about-us\)](#)

[Stuck in another language? Click here to go back to english \(?language=en\)](#)

Contact Us: admin@egap.org

MY ACCOUNT

[Log in \(/user\)](#)



- English (/methods-guides/10-things-you-need-know-about-causal-inference)
- Lang : English  Français (/fr/methods-guides/10-things-you-need-know-about-causal-inference)
- हिन्दी (/hi/methods-guides/10-things-you-need-know-about-causal-inference)
 - Português (/pt-br/methods-guides/10-things-you-need-know-about-causal-inference)
 - Español (/es/methods-guides/10-things-you-need-know-about-causal-inference)



10 Things to Know About Causal Inference

- Abstract
- 1 A causal claim is a statement about what didn't happen.
- 2 There is a fundamental problem of causal inference.
- 3. You can estimate average causal effects even if you cannot observe any individual causal effects.
- 4. If you know that, on average, A causes B and B causes C , this does not mean that you know that on average A causes C .
- 5. Causes are non rival.
- 6. It's easier to learn about the "effects of causes" than to learn about the "causes of effects."
- 7. Correlation is not causation.
- 8 X can cause Y even if X is not a necessary condition or a sufficient condition for Y .
- 9 Estimating average causal effects does not require that treatment and control groups are identical.
- 10 There is no causation without manipulation

Abstract

The philosopher David Lewis described causation as "something that makes a difference, and the difference it makes must be a difference from what would have happened without it."¹ This is more or less the interpretation given to causality by most experimentalists. It is a simple definition but it has many implications that can trip you up. Here are ten ideas implied by this notion of causality that matter for research strategies.²

1 A causal claim is a statement about what didn't happen.

For most experimentalists, the statement " X caused Y " means that Y is present but Y would not have been present if X were not present. The definition requires a notion about what could have happened (but didn't). This is the "counterfactual" (or sometimes "difference making") approach to causality and it can be distinguished from the "production" approach (which focuses on the idea of a causal connection between A and B). Under this approach there is no notion that just because X caused Y that X is the main reason or the only reason why Y happened.

Technical Note: Statisticians employ the "potential outcomes" framework to describe these counterfactual relations. In this framework we let $Y_i(1)$ denote the outcome for unit i that would be observed in condition 1 (e.g. treatment) and $Y_i(0)$ the outcome that would be observed, all else held constant, in condition 0 (e.g. control). The causal effect of X on Y is $Y_i(1) - Y_i(0)$. A treatment effect (or individual causal effect) is $Y_i(1) - Y_i(0)$.

The causal effect is then $\tau_i = Y_i(1) - Y_i(0)$. A treatment has a (positive or negative) causal effect on Y if $Y_i(1) \neq Y_i(0)$.³

2 There is a fundamental problem of causal inference.

If causal effects are statements about the difference between what happened and what could have happened, then causal effects cannot be measured. That's bad news. Prospectively, you can arrange things so you can see either what happens if someone gets a treatment or what happens if they do not get a treatment, but you cannot see both of these things and so you cannot see the difference between these two things. This is often called the "fundamental problem of causal inference."

3. You can estimate average causal effects even if you cannot observe any individual causal effects.

The fundamental problem notwithstanding, even if you cannot observe whether X causes Y in any given case, it can still be possible to figure out if X causes Y on average. The key insight here is that the average causal effect is the same as the difference between the average outcome for all units were they in the control condition and the average potential outcome for all units were they in the treatment condition. Many strategies for causal identification (see 10 Strategies for Figuring Out If X Caused Y (<http://egap.org/methods-guides/10-strategies-figuring-out-if-x-caused-y>)) focus on ways to figure out these average potential outcomes.

Technical Note: The key technical insight is that the difference of averages is the same as the average of differences. That is, using the "expectations operator," $\mathbb{E}(\tau_i) = \mathbb{E}(Y_i(1) - Y_i(0)) = \mathbb{E}(Y_i(1)) - \mathbb{E}(Y_i(0))$. The terms inside the expectations operator in the second quantity cannot be estimated, but the terms inside the expectations operators in the third quantity can be.⁴ See illustration here (<https://raw.githubusercontent.com/egap/methods-guides/master/causal-inference/PO.jpg>).

4. If you know that, on average, A causes B and B causes C , this does not mean that you know that on average A causes C .

You might expect that if A causes B and B causes C that therefore A causes C .⁵ But there is no reason to believe that average causal relations are transitive in this way. To see why, imagine A caused B for men but not women and B caused C for women but not men. Then on average A causes B and B causes C but there may still be no one for whom A causes C through B .

5. Causes are non rival.

Even if we focus uniquely on the effect of a single cause, X , on an outcome Y , we generally do not expect that is ever a single cause of Y .⁶ What's more, if you add up the causal effects of different causes there is no reason to expect them to add up to 100% so there is not much point trying to "apportion" outcomes to different causal factors. In other words causes are not rival. The National Rifle Association argues for example that guns don't kill people, people kill people. That statement does not make much sense in the counterfactual framework. Take away guns and you have no deaths from gunshot wounds. So guns are a cause. Take away people and you also have no deaths from gunshot wounds, so people are also a cause, and these two factors are simultaneously causes of the same outcomes.

6. It's easier to learn about the "effects of causes" than to learn about the "causes of effects."

Though it might sound like two ways of saying the same thing there is a difference between understanding what the effect of X on Y is (the "effects of a cause") and whether an outcome Y was *due* to cause X (the "cause of an effect").⁷ Imagine for example that X had a positive effect on Y for all men but a negative effect for all women. Then the average effect of X on Y would be 0. But for all cases with $Y = 1$, we see that $Y = 1$ *because* $X = 1$, and similarly for all cases with $Y = 0$, we see that $Y = 0$ *because* $X = 1$. Experimentation can get an exact answer to the first question, but generally it is not possible to get an exact answer to the second question.⁸

7. Correlation is not causation.

A correlation between X and Y is a statement about relations between actual outcomes in the world, not about the relation between actual outcomes and counterfactual outcomes. So statements about causes and correlations don't have much to do with each other. Positive correlations can be consistent with positive causal effects, no causal effects, or even negative causal effects. For example taking cough medication is positively correlated with coughing but hopefully has a negative causal effect on coughing.

Technical Note: Let D_i be an indicator that reports whether unit i has received a treatment or not. Then the difference in average outcomes between those that receive a treatment and those that do not can be written as $\frac{\sum_i D_i \times Y_i(1)}{\sum_i D_i} - \frac{\sum_i (1-D_i) \times Y_i(0)}{\sum_i (1-D_i)}$. This may or may not be a good estimate of difference in average potential outcomes for everyone. What matters is whether $\frac{\sum_i D_i \times Y_i(1)}{\sum_i D_i}$ is a good estimate of $\frac{\sum_i 1 \times Y_i(1)}{\sum_i 1}$ and whether $\frac{\sum_i (1-D_i) \times Y_i(0)}{\sum_i (1-D_i)}$ is a good estimate of $\frac{\sum_i 1 \times Y_i(0)}{\sum_i 1}$. This might be the case if those in treatment are a representative sample of the population, but otherwise there is no reason to expect that it would be.

8 X can cause Y even if X is not a necessary condition or a sufficient condition for Y .

We often talk about causal relations in deterministic terms. Even the Lewis quote at the top of this page seems to suggest a deterministic relation between causes and effects. Sometimes these are thought to entail necessary relations (for Y to occur X has to happen); sometimes they seem to entail sufficient relations (if X occurs then Y occurs). But once we are talking about multiple units there are at least two ways in which we can think of X causing Y even if X is not a necessary or a sufficient condition for Y (in fact some might think of these as being the same answer, given twice). One is to reinterpret everything in probabilistic terms: by X causes Y we simply mean that the probability of Y is higher when X is present. Another is to allow for contingencies — for example perhaps X causes Y if condition Z is present, but not otherwise.⁹

9 Estimating average causal effects does not require that treatment and control groups are identical.

People sometimes worry in experimental and other research designs that treatment groups and control groups look different. Very often experimental approaches are justified on the grounds that random assignment helps make sure that treatment and control groups are identical, "in expectation." But of course they might not be identical "in realization" (that is, in fact). Sometimes people even conduct statistical tests to see if the groups are identical. In fact, it is not possible to test for identicality in this way.¹⁰

fact in most applications they will never be identical in realization.¹⁰

The good news is that the argument for why differences in outcomes in randomly assigned treatment and control groups capture treatment effects does not rely on treatment and control groups being similar in their observed characteristics. In the absence of random assignment, treatment and control groups may look identical, but that in itself is no guarantee that they would act in the same ways, because they may differ in unmeasured ways. Conversely, a randomly assigned control group might look very different from a treatment group, but that does not take away from the fact that the average outcome in the control group gives an unbiased estimate of the average potential outcome $\mathbb{E}(Y_i(0))$ in the population.

10 There is no causation without manipulation

The definition of causal relations described above requires one to be able to think through how things might look in different conditions. How would things look if one party is elected compared to outcomes if another party is? But everyday causal statements often fall short of this requirement in one of two ways. First some statements do not specify clear counterfactual conditions. For example the claim that “the recession was caused by Wall Street” does not admit of an obvious counterfactual— are we to consider whether there would have been a recession if Wall Street did not exist? Or is the statement really a statement about particular actions that Wall Street could have taken but did not. If so, which actions? The validity of these statements is a bit hard to assess, and can depend on which counterfactual conditions are implied by the statement. Perhaps a bigger problem arises when counterfactual conditions cannot even be imagined. For example the claim that Peter got the job because he is Peter implies a consideration of what would have happened if Peter was not Peter (or for another example, the claim that Peter got the job because he is a man requires considering Peter as a woman). The problem is that the counterfactual implies a change not just in the condition facing an individual but in the individual themselves. To avoid such problems some statisticians urge a restriction of causal claims to treatments that can conceivably (not necessarily practically) be manipulated.¹¹ While we might have difficulties with the claim that Peter got the job because he was a man, we have no such difficulties with the claim that Peter got the job because the hiring agency thought he was a man.

-
1. Lewis, David. “Causation.” *The journal of philosophy* (1973): 556-567.↵
 2. Originating author: Macartan Humphreys. Minor revisions: Winston Lin and Donald P. Green, 24 Jun 2016. Revisions MH 6 Jan 2020. The guide is a live document and subject to updating by EGAP members at any time; contributors listed are not responsible for subsequent edits.↵
 3. Holland, Paul W. “Statistics and causal inference.” *Journal of the American Statistical Association* 81.396 (1986): 945-960.↵
 4. Holland, Paul W. “Statistics and causal inference.” *Journal of the American Statistical Association* 81.396 (1986): 945-960.↵
 5. Interpret “ A causes B , on average” as “the average effect of A on B is positive.”↵
 6. In some accounts this has been called the “Problem of Profligate Causes”.↵
 7. Some reinterpret the “causes of effects” question to mean: what are the causes that have effects on outcomes. See Andrew Gelman and Guido Imbens, “Why ask why? Forward causal inference and reverse causal questions”, NBER Working Paper No. 19614 (Nov. 2013).↵
 8. See, for example, Tian, J., Pearl, J. 2000. “Probabilities of Causation: Bounds and Identification.” *Annals of Mathematics and Artificial Intelligence* 28:287–313.↵
 9. Following Mackie, sometimes the idea of “INUS” conditions is invoked to capture the dependency of causes on other causes. Under this account a cause may be an Insufficient but Necessary part of a condition which is itself Unnecessary but Sufficient. For example dialing a phone number is a cause of contacting someone since

having a connection and dialing a number is sufficient (S) for making a phone call, whereas dialing alone without a connection alone would not be enough (I), nor would having a connection (N). There are of course other ways to contact someone without making phone calls (U). Mackie, John L. "The cement of the universe." London: Oxford Uni (1974).↵

10. For this reason *t*-tests to check whether "randomization worked" do not make much sense, at least if you believe that a randomized procedure was followed. If there are doubts about whether a randomized procedure was correctly implemented these tests can be used to test the hypothesis that the data was indeed generated by a randomized procedure.↵
11. Holland, Paul W. "Statistics and causal inference." Journal of the American Statistical Association 81.396 (1986): 945-960.↵

More Methods Guides

[- View Another Methods Guide -](#)

Go

Methods Guide Authors

Macartan Humphreys (</content/macartan-humphreys>)

