Separating causal questions from statistics with application to dietary interventions

SPER Advanced Methods Workshop

April 28th, 2026

Structure of workshop

Part 1 (Jessica): Some foundations

- Choosing and defining a causal effect in which to ground causal inference using counterfactuals
- Introduce the generalized g-formula as a foundation for estimating many different types of causal effects under causal assumptions.

Part 2 (Lan): Different methods for estimating a generalized g-formula that follow from different statistical assumptions (g-computation, IPW and DR methods)

Part 3 (Soren): A practical application of these concepts to estimate effects of dietary interventions in Project Viva data.

Causal inference versus statistical inference?

What is the difference?

Statistics and statistical inference

- We can understand a *statistic a*s any computation applied to a data set on a sample of individuals from a study population of interest
- *Statistical inference* is the process of using a statistic to learn about some ***factual (measureable)*** feature of that population
- We can refer to the factual feature of the population we wish to use a statistic to learn about is the ***statistical target***.

Each dart represents a statistic computed from data in one sample

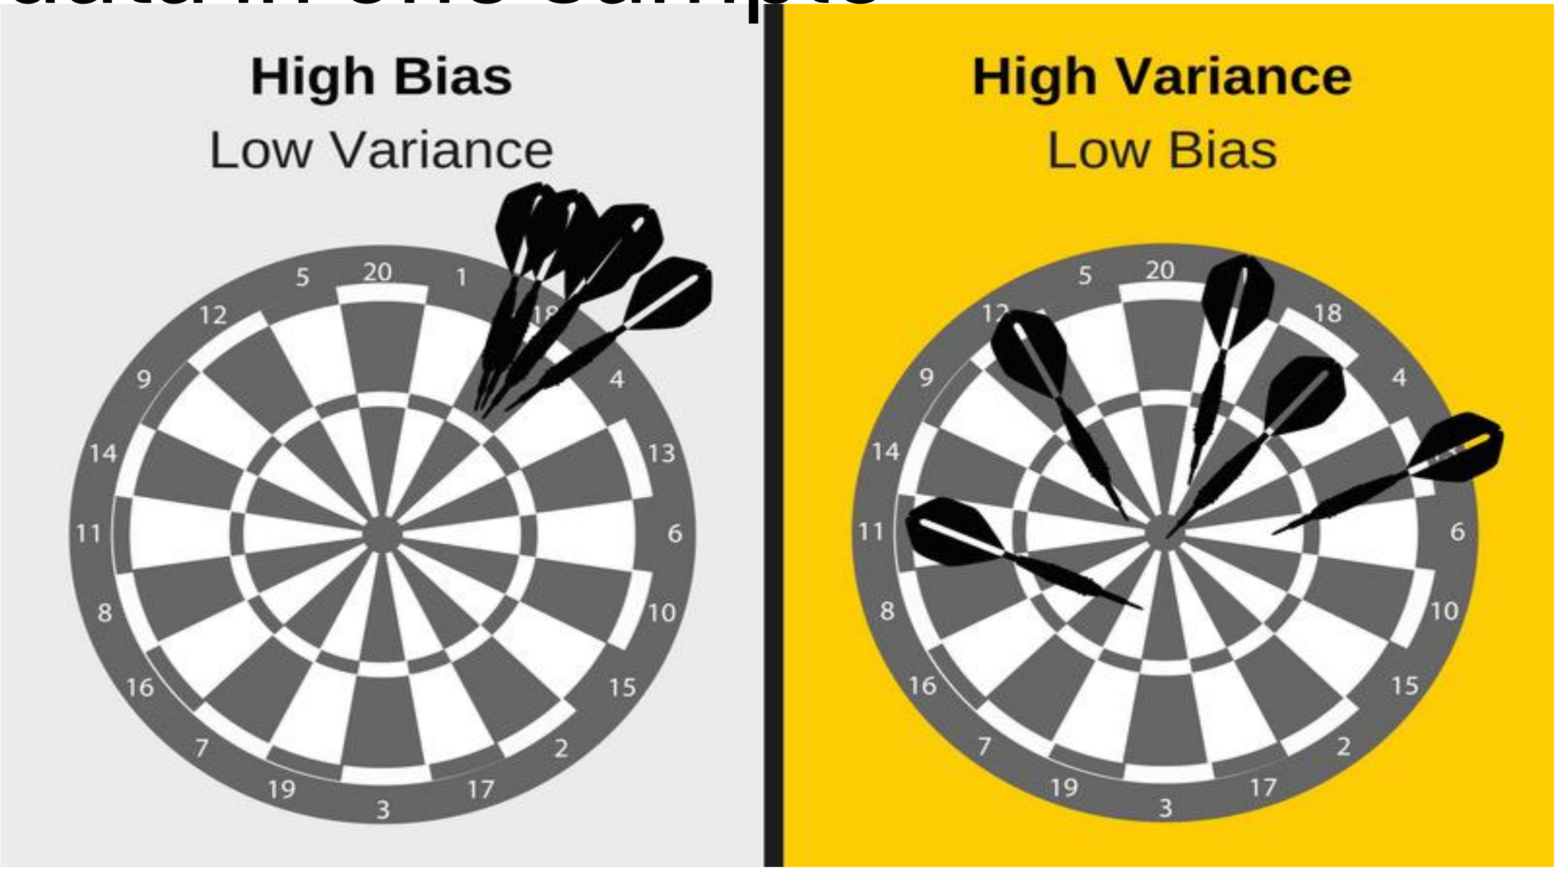

The center represents the measureable population feature of interest: the "statistical target"

Causal inference?

Causal inference can be understood as the process of learning about a ***causal effect*** in a study population of interest

- Causal effect: inherently defined by **counterfactual** features of that population

A causal effect therefore does not constitute a statistical target (it is not factual/measureable).

- Requires knowledge of what would happen to the same population under two different scenarios (only one of those can possibly be factual).

To use statistics for the task of causal inference

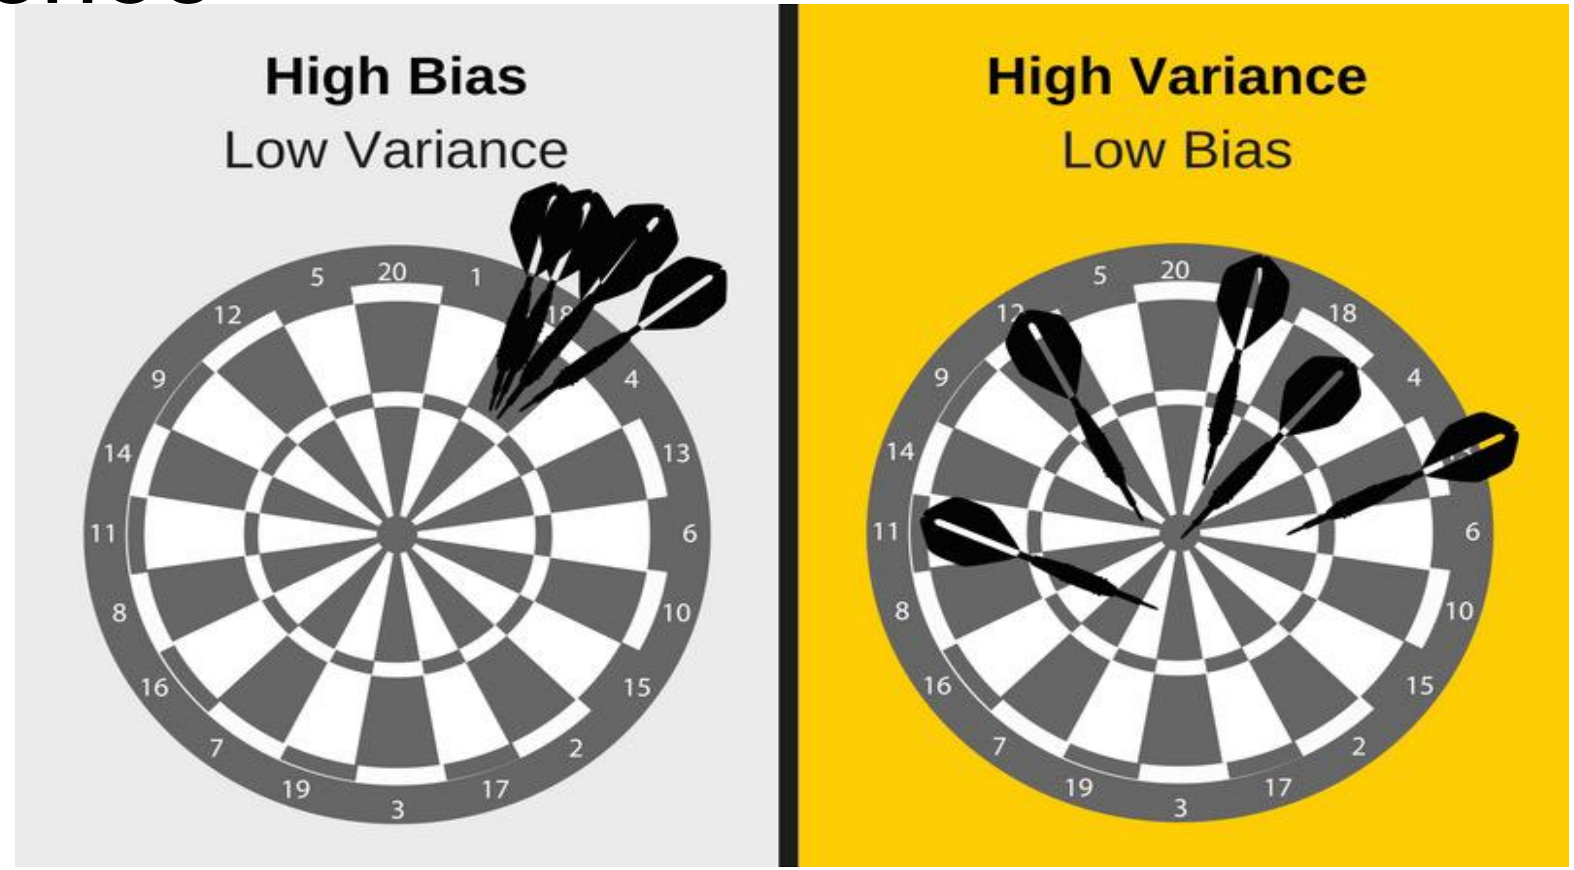

We need to be able to equate our causal effect to a statistical target.

Assumptions

To make that link between a causal effect and a statistical target – and ultimately a choice of statistic - we need assumptions.

Assumptions

Causal effect we want

Statistic we apply to the data we have

Two types of assumptions

1. “Causal” identifying assumptions: these are what allow us to link our causal effect (a counterfactual population feature) to a statistical target (a factual/measureable population feature).
    - Ex. “no unmeasured confounding”
2. “Statistical” assumptions: these determine properties of different statistics for that statistical target.
    - Ex. We can correctly specify a parametric outcome regression model.

Example

Consider a study where the following variables are measured on a sample from the study population of interest (relative to some baseline time).

- A: weekly fast food consumption over some time period (e.g. the subsequent year). Simplifying assumption: this is “time fixed” (no change in this behavior over this period)
  - $A=0$: no consumption; $A=1$: 1-2x per week; $A=2$: 3+ x per week;
- Y: health measure of interest (Soren example: HOMA-IR score)
- L: a set of covariates measured at baseline.

Causal effect to ground a choice of statistic?

Suppose we consider interest in a causal effect defined as

The mean difference in the outcome had, over that same year period

- everyone consumed no fast food

versus had instead

- everyone consumed at least 3 servings per week

Because our causal effect refers to an intervention that would manipulate fast food consumption, we will call the variable A our "treatment".

A more formal definition that could be linked to a statistical target

Let Y^a be an individual's value of the outcome had, possibly contrary to fact, they adhered to an intervention that required fast food consumption take particular fixed value a.

In turn, we could define the target causal effect in this case in terms of the counterfactual mean (average) difference

$$E(Y^{a=2}) - E(Y^{a=0})$$

where each $E(Y^a)$ is the mean of the outcome we would see if everyone in the population had consumed that level a.

Stage 1 assumptions

Given interest in this causal effect, we need to think about a set of *causal identifying assumptions* that bring us to a choice of statistical target.

Causal effect we want

Statistic we apply to the data we have

This will tell us what the center of the dart board is!

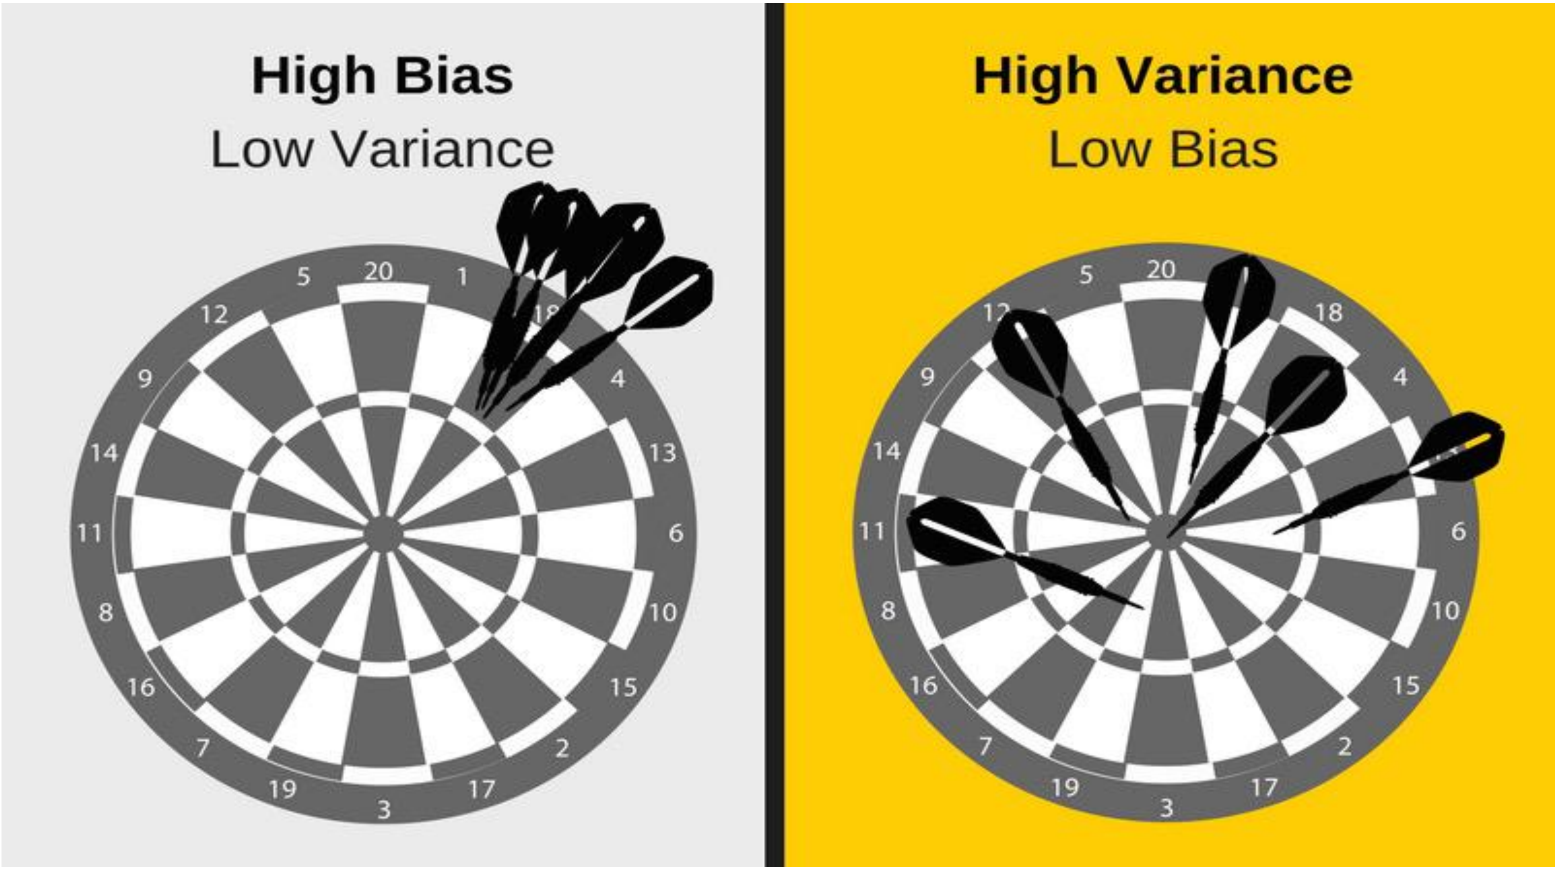

One set of assumptions premised on “no unmeasured confounding”

Suppose we could reasonably assume that, within joint levels of the measured covariates L, fast food consumption was “effectively randomized”.

- This would be true by design if an individual’s level of treatment A was physically assigned by a fair coin toss within levels of L in the study.
- In this case, we can expect that Y^a and A are independent conditional on L.

This independence is referred to as **conditional exchangeability**

It formalizes the assumption of “no unmeasured confounding”. Not guaranteed without physical randomization of A conditional on L.

A familiar way to reason about conditional exchangeability: Causal DAGs

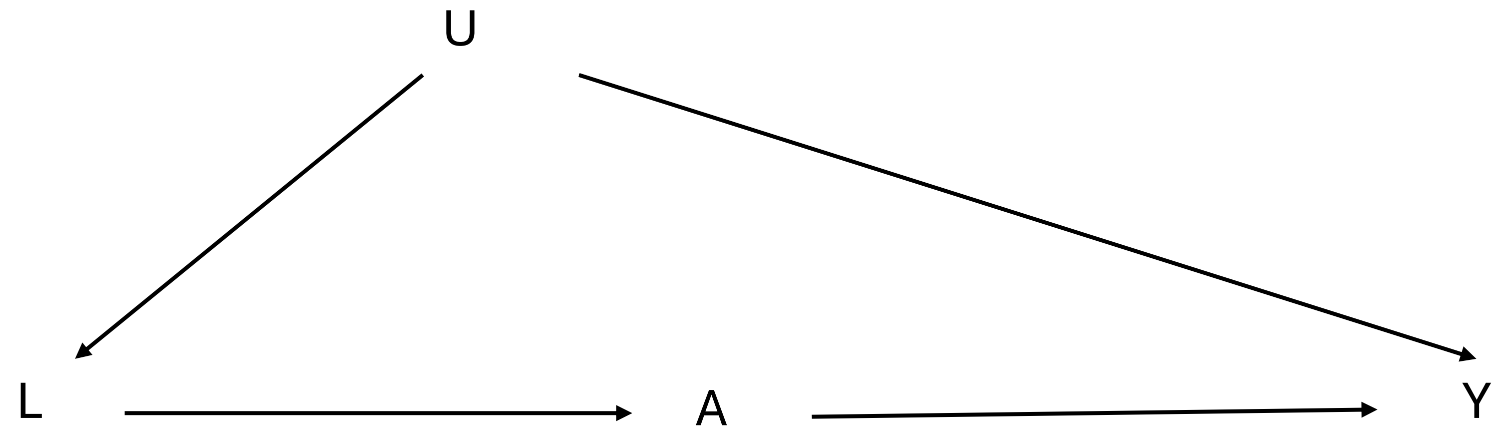

- No “unblocked backdoor paths” between treatment (A) and outcome (Y) conditional on measured covariates (L)

A less familiar way: SWIGs

- A causal DAG is an **implicit** representation of a *counterfactual* causal model (there are no counterfactuals on it!)
- It is not as obvious to understand on a causal DAG why conditional exchangeability would hold or fail.
- Single World Intervention Graphs (SWIGs) are an alternative way to communicate a causal model that explicitly show counterfactuals.
- *Single World* refers to the specific “counterfactual world” in which an intervention of interest is implemented in all individuals.

One way to construct a SWIG is by starting with a causal DAG and applying certain *intervention-specific* transformations.

Construction of a SWIG from a DAG

1. Split each treatment node:
   - the right side of the split represents the intervention treatment value,
   - the left the "natural treatment value" (NTV).
   - For "time-fixed" treatment the NTV is the observed value of treatment (A).
2. All arrows on the DAG going out of A now go out of the intervention value on the SWIG. All arrows going into A on the DAG stay into A on the SWIG.
3. All variables "affected by" the intervention treatment value are counterfactually labeled.
4. If, under the intervention, the intervention treatment value would be determined by a past variable, an arrow is added from that variable into the intervention treatment value. Not relevant for the strict intervention "everybody gets level a".

Example SWIG under an intervention where everyone's treatment is strictly set to fixed a

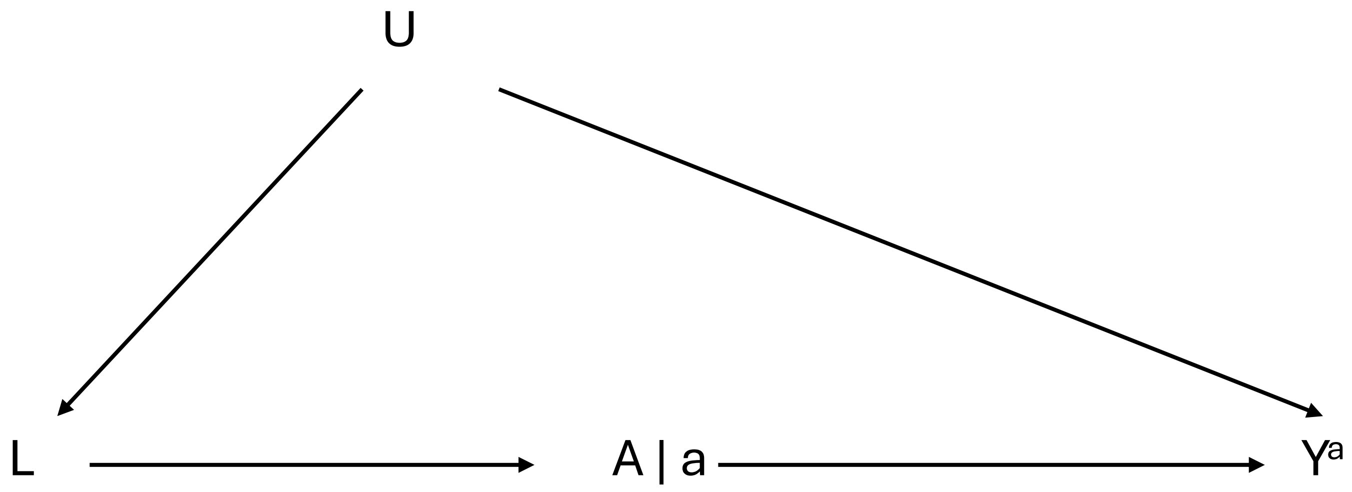

- Y^a independent of A is evaluated by checking for no unblocked backdoor paths between Y^a and A on the SWIG.

Other causal identifying assumptions “tied to” conditional exchangeability

Two other causal identifying assumptions are (minimally) needed when we rely on a conditional exchangeability (no unmeasured confounding) assumption to connect a causal effect to a statistical target

- Consistency
- Positivity

Consistency

- When the intervention of interest is “set treatment to fixed a” consistency states: If A=a then $Y=Y^a$

Constructing a SWIG from a causal DAG can be understood as inherently communicating a “belief” in consistency.

Positivity

Positivity is the assumption that

- Conditional on any possible level of the ”measured confounders” (L), *any treatment level that is* ***possible to see under the intervention of interest is also possible to see in the data*** *(in factual reality)*.
- When the intervention of interest is “set treatment to fixed a” positivity requires: $P(A=a|L=l)>0$ for any level of l that is possible in the data.
  - It is possible to see people who factually got level a within every possible combination of the measured confounders.

The statistical target: the g-formula

When these three assumptions hold for a strict intervention requiring that "everyone receive the same treatment level a", we can write $E(Y^a)$ in terms of a particular function of the data: the g-formula.

$$\Sigma_l \{E[Y|A=a, L=l\,] \times \Pr[L=l\,]$$

A contrast in thisg-formula for the two levels of a is now our statistical target: we got to it from starting with a particular causal effect of interest and invoking particular causal identifying assumptions.

The "right" statistical target?

This is only the "right" statistical target (the thing we should go after with some statistic) provided

1. The identifying assumptions we invoked are reasonable
2. This causal effect is really the question we want answered!

Understanding the g-formula

$$\Sigma_l \; E[Y|A=a,L=l\,] \times \Pr[L=l\,]$$

- This is a weighted sum of the conditional outcome mean $E[Y|A=a,L=l\,]$ over all levels of the measured confounders L.

- $\Pr[L=l]$ is the weight in the weighted sum– the proportion of people with level $L=l$ in the study population

- The symbol Σ_l reads "the sum over all possible levels of the measured confounders L"

From causal effect to choice of statistic

We have just outlined a thought process from a choice of causal effect to a statistical target. This is only the first half of the journey. The second half requires a thought process from statistical target to statistic (this will depend on statistical assumptions when there are many levels of L).

Causal effect we want

Statistic we apply to the data we have

Reflecting on this choice of causal effect?

Is this choice of causal effect what we want? Is it *useful*?

- We considered a comparison of outcomes under a scenario where everyone would eat high amounts of fast food compared to none.
- Why would we care about an intervention that forces *everyone* to eat fast food 3 or more times per week? Even individuals who, in fact, ate no fast food!
- While we might care about an intervention that forces *everyone* to eat no fast food, is it implementable/practical?

The caveats of strict/non-pragmatic causal questions: utility

The causal inference literature and associated pedagogy has been historically limited to consideration of only strict intervention rules like the example we just considered

- Intervention rules that require everyone in a target study population to take the same value of “treatment” under intervention

Such rules are often divorced from any real policy/clinical decision which, in reality, will need to take into account information specific to individuals.

The caveats of strict/non-pragmatic causal questions: identifiability

Often lack of "utility" goes hand in hand with lack of "identifiability", particularly regarding the positivity assumption

Recall what positivity requires:

- *positivity requires that, conditional on any level of the measured confounders possible in the data, any treatment level that is possible to see under the specified intervention is also possible to see in the data*.

When interventions are defined unrealistically, there will often be treatment patterns possible under intervention that are impossible (or close to impossible) in reality.

Fortunately…

- We are not restricted to grounding our causal inferences in strict (non-pragmatic/unrealistic) questions.

- We need not force ourselves into considering effects that contrast counterfactual means of the form $E(Y^a)$ for two choices of a when that does not make sense from a clinical or policy perspective.

Definition of a pragmatic intervention on fast-food

- Contemporary dietary counselling:
  - ≠ one-size-fits-all recommendations
  - ≠ unidirectional, knowledge-based session
- Intervention level of treatment is determined based on <u>where</u> the client/patient is at (their “natural” consumption level)

Consumes less than 1 d/week → 🙌 You continue!

Consumes 1-2 d/week → 💪 Reduce to less than 1 day/week

Consumes at least 3 d/week → 👍 Reduce to 1-2 day/week

Cullen, Baranowski, Smith. J Am Diet Assoc 2001

Natural treatment value intervention

- This is an example of an intervention that depends on an individual's "natural treatment value", the observed A in our case.
- In this example, the intervention rule is defined as
  - If A=0, then, under intervention, treatment should take the value 0.
  - If A>0, then, under intervention, treatment should take the value A-1.

  Under this intervention rule, everyone will NOT have the same treatment value. It will depend on what they would do "naturally".
- We might define an effect that contrasts outcome means under this pragmatic intervention versus no intervention (always receive your natural treatment value).
- This would capture the effect of following a pragmatic guideline for reducing fast food consumption compared to no change.

A different starting point

This is a different causal effect and (maybe not surprisingly), this is going to lead us to a slightly different destination.

Assumptions

Causal effect we want

Statistic we apply to the data we have

Generalized causal inference

This is one example of what we can call a "generalized" intervention rule: an intervention rule where treatment under intervention depends on some part of the individual's "measured past" (A and/or L).

We can denote the outcome mean under such a general intervention rule $E(Y^g)$.

The more familiar $E(Y^a)$ is just a special case of a g where there is no dependence of the intervention treatment value on A or on L.

Causal identifying assumptions

- We can analogously identify a counterfactual mean under a generalized intervention $E(Y^g)$ under conditional exchangeability, consistency, and positivity.

- These assumptions are now defined with respect to whatever g we consider.

- Conditional exchangeability: Y^g independent of A conditional on L

Evaluating conditional exchangeability: g-specific SWIG, NTV intervention example

- A^{g+} is the intervention treatment level. The dashed red arrow is added for a choice of g where the intervention level depends on A.

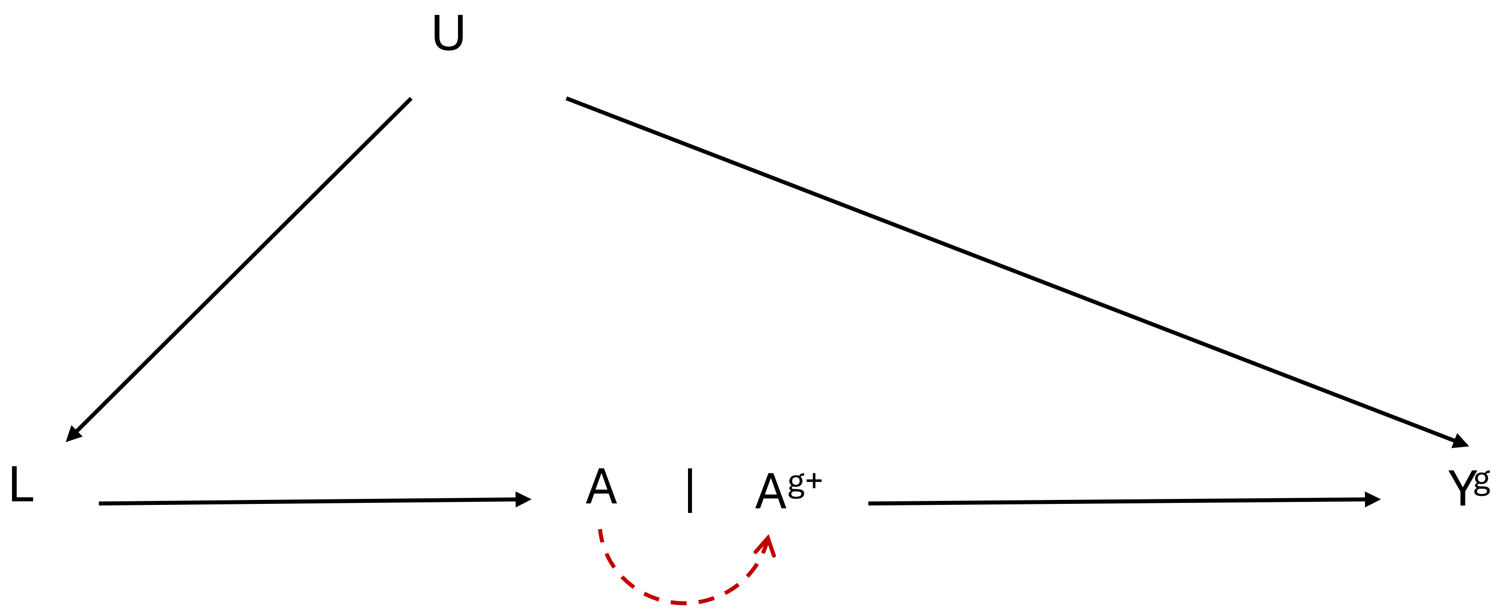

- As before, conclude Y^g independent of A conditional on L if no unblocked backdoor paths between Y^g and A given only L

Positivity

Positivity is the assumption that

- Conditional on any possible level of the ”measured confounders” (L), *any treatment level that is* ***possible to see under the intervention of interest is also possible to see in the data*** *(in factual reality).*

An example of how the choice of intervention impacts positivity

For our strict effect, that considers an intervention "everyone must eat 3+ servings per week", positivity requires $P(A=2|L=l)>0$ for any level of l that is possible in the data.

This would fail if there is any level of L where we don't have any people with this observed treatment level.

Our pragmatic rule doesn't require anyone to have 3+ servings because under the rule, no one would be forced to take this value.

- Those who would naturally eat no fast food would eat no fast food under intervention
- Those who would naturally eat 1-2 servings would also eat no fast food under intervention.
- Those naturally eating 3+ would eat 1-2 servings under intervention.

Generalized statistical target

For this general intervention g, we can identify $E[Y^g]$ under conditional exchangeability, consistency, and positivity assumptions specific to that intervention via a generalized version of the g-formula

$$\Sigma_l \Sigma_a E[Y|A=a,L=l\,] \times \Pr[A^{g+}=a|L=l\,] \times \Pr[L=l\,]$$

This is a weighted sum just like before. But it differs from the "strict" g-formula in two ways highlighted in red.

Understanding the ”generalized” g-formula

$$\Sigma_l \Sigma_a E[Y|A=a,L=l] \times \Pr[A^{g+}=a|L=l] \times \Pr[L=l]$$

- The conditional outcome mean is now additionally weighted by the chance of having a particular level of treatment under g conditional on confounder level L=l ($\Pr[A^{g+}=a|L=l]$).
- The sum is now also over all possible intervention treatment levels Σ_a
- We have this second sum because we allow now that under g there can be more than one possible intervention treatment level.

Special cases of the generalized g-formula

We can gain additional intuition on

$$\Sigma_l \Sigma_a \mathrm{E}[Y|A=a,L=l\,] \times \Pr[A^{g^+}=a|L=l\,] \times \Pr[L=l\,]$$

by considering some special cases of g

Special case: when g is selected as strict strategy

Suppose we choose the intervention to be a strict strategy: force everyone to have treatment level a=2.

$$\Sigma_l \Sigma_a \mathrm{E}[Y|A=a,L=l\ ] \times \Pr[A^{g+}=a|L=l\ ] \times \Pr[L=l\ ]$$

In this special case the chance of having treatment level a=2 (regardless of L) under g is 1! And the chance of having any other level (regardless of L) is 0!

Translation: $\Pr[A^{g+}=a|L=l\ ]=1$ for $a=2$ and $\Pr[A^{g+}=a|L=l\ ]=0$ for any other value of a

Reduces to the strict g-formula: $\Sigma_l \mathrm{E}[Y|A=2,L=l\ ] \times \Pr[L=l\ ]$

Special case: when g is selected as no intervention

Suppose we choose the intervention to be “no intervention” (always set A^{g+} to A)

$$\Sigma_l \Sigma_a E[Y|A=a,L=l\,] \times \Pr[A^{g+}=a|L=l\,] \times \Pr[L=l\,]$$

In this special case the chance of having any given treatment level a within a given level of L under intervention will be the same as it is in reality (the factual world).

Translation: $\Pr[A^{g+}=a|L=l\,] = \Pr[A=a|L=l\,]$

The generalized g-formula in this case reduces simply to the factual outcome mean E[Y].

What is coming up next

Part 2 (Next): Lan will walk us through different statistics for generalized g-formula (both the special case of a strict g and our pragmatic natural value intervention).

For the same g-formula (statistical target), different statistics will follow from different statistical assumptions.

Part 3: Soren will walk us through an application of these ideas in a practical example along with R code in a simulated data set.

Some resources (application papers)

- Lajous et al. Changes in fish consumption in midlife and the risk of coronary heart disease in men and women. AJE 2013.
- Sewak et al. Causal effects of stochastic PrEP interventions on HIV incidence among men who have sex with men. AJE 2024.
- Rudolph KE et al. When effects cannot be estimated: redefining estimands to understand effects of naloxone access laws. Epidemiology, 2023.
- Harnois-Leblanc S. et al. Separating causal questions from statistics with application to dietary interventions in the Project Viva cohort offspring. In preparation.

Some technical resources (conceptual/identification)

- Richardson TS, Robins JM. Single World Intervention Graphs (SWIGs): A unification of the counterfactual and graphical approaches to causality. April 2013 CSSS Working Paper #128.
- Young et al. Identification, estimation and approximation of risk under interventions that depend on the natural value of treatment using observational data. Epi Methods. 2014
- Sarvet AL, Stensrud MJ. The natural value of treatment and its importance for causal inference. Annual Review of Statistics and Its Application. 2026.

More technical resources (statistics)

- Haneuse S, Rotnitzky A. Estimation of the effect of interventions that modify the received treatment. Stat in Med. 2013.
- Diaz et al. Nonparametric causal effects based on longitudinal modified treatment policies. JASA Theory and Methods. 2023
- Wen et al. Intervention treatment distributions that depend on the observed treatment process and model double robustness in causal survival analysis. SMMR. 2024.