Inverse probability weighting and the parametric g-formula for estimating effects of time-varying dynamic treatment strategies in observational studies

Jessica G. Young

Department of Population Medicine
Harvard Medical School & Harvard Pilgrim Health Care Institute

April 24, 2019

Workshops so far

- Yi: reviewed idea of "emulating a target trial with observational data" (motivating examples: time-varying EPO strategies)
  - brief intro to inverse probability weighting (IPW) and parametric g-formula
- Maya: Formal framework for causal effects of point treatments as well as (deterministic) static time-varying strategies; introduction to dynamic strategies
  - TMLE estimator of these effects coupled with Superlearner

Today

Deeper dive into:

- Defining causal effects of time-varying treatment strategies on an outcome mean (focus on risk of a failure event/survival case)
  - Will cover general *deterministic* treatment strategies
  - Can be classified as either *static* or *dynamic* – will review examples
- Assumptions needed for identifying these effects with observational data from a longitudinal study
- How choice of causal effect + assumptions + available data leads us to estimators
  - Inverse probability weighted estimators
  - Parametric g-formula
- And tradeoffs of methods – why ever use anything but TMLE with superlearner?

Motivating example

Over several decades, various questions were posed about

- *causal effects of following different time-varying antiretroviral therapy (ART) initiation strategies*

on long term all-cause mortality risk in ART-naive HIV-infected populations. Strategies=interventions= rules.

Effects of deterministic static interventions

We might consider the causal effect on 5-year all-cause mortality risk in this study population of the interventions:

- Initiate ART at baseline versus
- Never initiate ART at any time within 5-years of baseline.

These strategies/rules/interventions can be classified as *deterministic static interventions*:

- treatment assignment at any time during follow-up is fully determined at baseline.

Effects of deterministic dynamic interventions

Alternatively, we might consider the causal effect on 5-year all-cause mortality risk in this study population of the interventions:

- If CD4 cell count first drops below 450 at time k then initiate treatment at k; otherwise do not initiate at time k versus
- If CD4 cell count first drops below 350 at time k then initiate treatment at k; otherwise do not initiate at time k

These strategies/rules/interventions can be classified as *deterministic dynamic interventions*:

- treatment assignment at a given time depends on individual risk factor history and is fully determined by that history.

Effects of random dynamic interventions

Yet a third option... we might consider the causal effect on 5-year all-cause mortality risk in this study population of the interventions:

- If CD4 cell count first drops below 450 at time k then initiate treatment at k with probability $2/3$; otherwise initiate at time k with probability $1/4$
- If CD4 cell count first drops below 350 at time k then initiate treatment at k with probability $2/3$; otherwise initiate at time k with probability $1/4$

These strategies/rules/interventions can be classified as *random dynamic interventions*

- treatment assignment at a given time depends on individual risk factor history but is *not* fully determined by that history.
- treatment is assigned at a given time by a weighted flip of a coin depending on risk factor history

Choosing a question?

So there are many possible ways to define the causal effect of time-varying treatment rules

- Because there are an infinite number of ways to define these rules

How do we choose?

- Will involve a tradeoff between the question we want to answer and the data we have available to answer it.

Ideal Data

To answer any of these questions, the ideal study would

1. Randomize individuals from the study population at baseline to one of the two strategies (however defined, static, dynamic, deterministic, random)
2. Force adherence to the rule required by the baseline randomization throughout the follow-up (5-years or death, whichever comes first)
3. Simply compare the difference in proportions of death from any cause at the end of 5-years in each arm
   - The number of deaths in arm 1 divided by number assigned to arm 1 at baseline vs. number of deaths in arm 2 divided by number assigned to arm 2 at baseline

Pretty easy. No fancy methods needed (e.g. two-sample t-test)

Real world data

- Rarely is such a study conducted
- Even if a trial is conducted, cannot ethically force adherence to any rule over time
- Trials not always feasible or timely so best available data may come from an observational cohort study (no baseline randomization)
- Whether a randomized or nonrandomized study, people may drop out such that outcome status is missing

HIV-CAUSAL Collaboration

- Includes several cohort studies from five European countries and the United States
- Cohorts assembled prospectively and based on data collected for clinical purposes within national health care systems with universal access to care

These data have been used to estimate causal effects of many types of ART initiation strategies

- When to start treatment based on CD4 cell count? Cain et al. (IPW) and Young et al. (parametric g-formula)

Definition of the Study population/Baseline

In these analyses, HIV-infected study population was defined as those who

- are at least 18 years old
- have never previously initiated antiretroviral therapy (ART)
- no history of CD4 cell count less than 500 cells/mm^3
- not pregnant
- have CD4 cell count and viral load (HIV RNA) measurements within last 6 months

Define the first time at which an individual meets all of these criteria *and* CD4 first drops into range 200-499 cells/mm^3 as "time 0" (baseline). Baseline may correspond to different calendar times for different people (here considered individuals meeting eligibility criteria between 1996 and 2009).

Observed data (what we measured)

Monthly measurements of the following for each subject, $k = 0, \ldots, K = 59$

- A_k: indicator of ART initiation by month k
- L_k: vector of mortality risk factors (e.g. most recent CD4 measurement, viral load, AIDS, whether a lab measurement taken in month k)
- Y_{k+1}: mortality status by the next month
- C_{k+1}: censoring by the next month (censored after 12 consecutive months with no lab measurement)

L_0 also includes time-fixed baseline characteristics (e.g. sex, race).
$Y_{K+1} \equiv Y_{60}$ (outcome status by 60 months).

- "history" of a variable denoted with overbars: e.g. $\overline{A}_K = (A_0, \ldots, A_K)$

Assume temporal order (C_k, Y_k, L_k, A_k) within each month k

Deterministic effects

With these questions and HIV-CAUSAL in mind, will consider a general framework for

1. Defining causal effects of time-varying deterministic strategies (static or dynamic)
2. Assumptions for identifying these effects in observational studies like HIV-CAUSAL
3. Function of observed data that identifies these effects under these assumptions– the g-formula
4. IPW and parametric g-formula estimators

Will briefly consider extensions to random dynamic strategies – and consider why we might choose them.

Counterfactual outcomes

Denote g as any deterministic time-varying ART initiation strategy

- Define Y_{60}^{g} as the indicator of whether an individual in the study population would die of any cause by 60-months post-baseline had, *possibly contrary to fact*, he adhered to the rule g
  - We say this is that individual's *counterfactual outcome* under the rule g
  - $Y_{60}^{g} = 1$ for an individual who would die of any cause within 60 months of his/her time 0 under g.
  - $Y_{60}^{g} = 0$ for an individual who would be alive under g by this time.

We are going to consider implicit in any definition of an intervention g: "eliminate censoring" (so interpretation of counterfactual actually depends on our definition of censoring)

- " ....and (somehow) force everyone to have a lab measure at least every 12 months".

Counterfactual definition of causal effect

For any two different deterministic ART initiation rules g_1 and g_2 define

- Individual causal effect: $Y_{60}^{g_1} - Y_{60}^{g_2}$
- Average causal effect: $\Pr[Y_{60}^{g_1} = 1] - \Pr[Y_{60}^{g_2} = 1]$ (causal risk difference)

E.g. g_1 versus g_2 could be "Initiate at baseline" versus "Never initiate" or "Initiate when CD4 first drops below x, otherwise don't" for $x = 450$ versus 350.

The Challenge

- Patients (and doctors) really want to know individual causal effects: clearly individual effects will require strong assumptions for identification in any study (even a trial). So we won't go there...
- The average causal effect is less of a reach – at least directly identified in ideal trial

In observational studies (not the ideal trial) we need assumptions that may (or may not) be reasonable.

Identifying Assumptions

Under assumptions, it is possible to link this observed longitudinal data to our counterfactual question (the average causal treatment effect of g_1 versus g_2):

1. Exchangeability
2. Positivity
3. Consistency

Exchangeability

For each k and both choices of g:

$$Y^g_{K+1}, \ldots, Y^g_{k+1} \coprod (A_k, C_{k+1}) | \overline{L}_k, \overline{A}_{k-1} = \overline{a}^g_{k-1}, \overline{Y}_k = \overline{C}_k = 0$$

In English:

- At each follow-up time k, an individual's observed treatment status (and censoring) is independent of his/her future counterfactual outcomes under the rule g conditional on his/her observed past history of treatment and covariates, this history being consistent with the intervention g and survival to k

This assumption is untestable because we don't observe outcomes under g for everyone. Holds by design in a sequentially randomized trial where A_k assigned by weighted coin based on past values of $\overline{L}_k, \overline{A}_{k-1}$ (and no one is censored). In general, not guaranteed to hold.

No unmeasured confounding

This conditional counterfactual independence assumption

$$Y^g_{K+1}, \ldots, Y^g_{k+1} \coprod (A_k, C_{k+1}) | \overline{L}_k, \overline{A}_{k-1} = \overline{a}^g_{k-1}, \overline{Y}_k = \overline{C}_k = 0$$

is equivalently called "No unmeasured confounding"

- We say in this case that $\overline{L}_k$ is the "measured confounder history through k" (e.g. measured CD4 history, viral load history, AIDS status through k, clinic visit history)

Causal diagrams to evaluate exchangeability

- Exchangeability cannot be assessed in the data
- It can be assessed under an assumption on how the data in the study were generated driven by subject matter knowledge (or simply by assumptions when knowledge is absent)
- Causal directed acyclic graphs (DAGs) can be used to represent these data generating assumptions
- Under these assumptions, exchangeability (or failure of this assumption) can be evaluated by certain manipulations of the DAG

Causal DAG representing sequential randomization

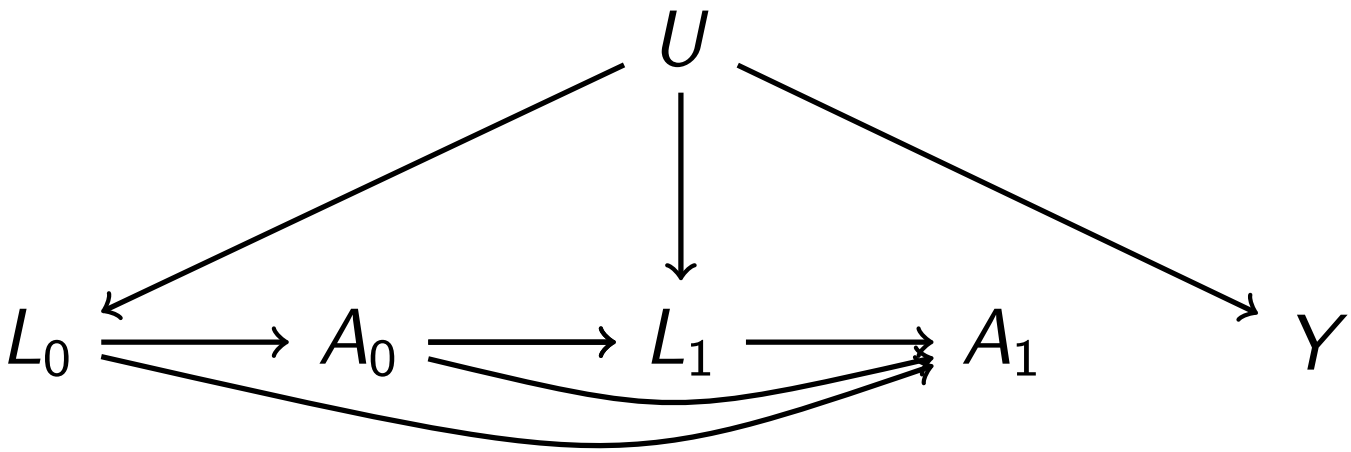

This causal DAG is consistent with data generating assumptions of a sequentially randomized trial where A_k is assigned by a weighted flip of a coin dependent only on $(\overline{L}_k, \overline{A}_{k-1})$ (and no censoring). Allows that there is an unmeasured common cause of mortality and measured confounders such as CD4 cell count (U could be genetic factors). If data arose from such a trial, design guarantees no arrow from U into A at any time.

Causal DAG representing HIV-CAUSAL data generating assumption

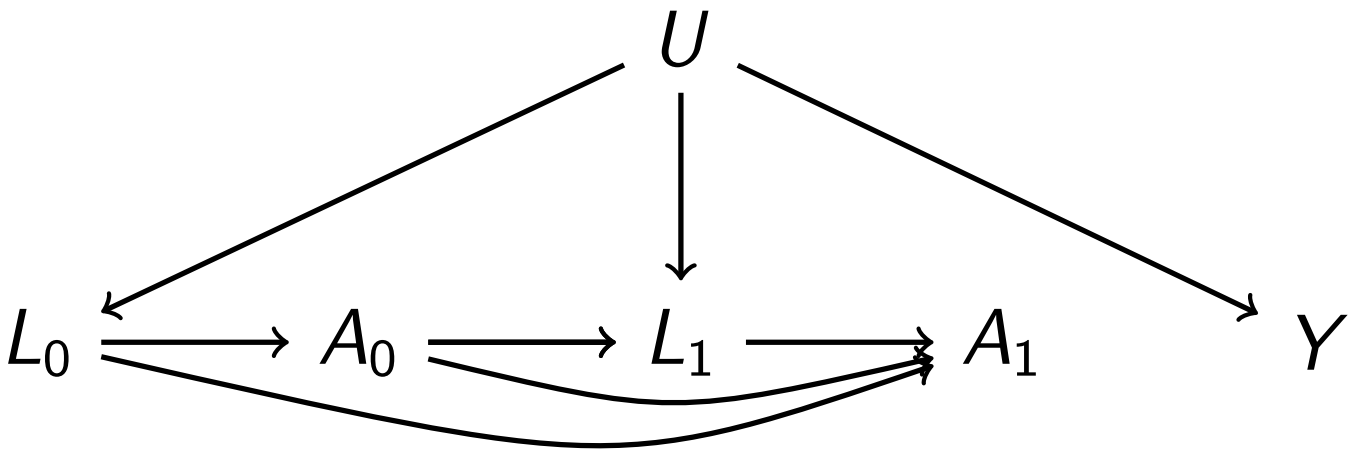

As we did not physically assign treatment initiation in HIV-CAUSAL, the absence of arrows from U into A at any time is an assumption. Reasonable assumption if U is genetics and doctor also knew nothing about genetics. If U is something doctor saw to determine treatment that we don't have in data set, need to draw an arrow.

General evaluation of exchangeability

In general, exchangeability for any strategy g guaranteed on a causal DAG where only arrows from $(\overline{L}_k, \overline{A}_{k-1})$ go into A_k at each time k.

- This is a special case, can define a more general approach to reading exchangeability off of a causal DAG by the absence of "unblocked backdoor paths" (Pearl)

Extra slides at end with details and example.

Positivity Assumption

General definition of positivity assumption for any choice of strategy g (static or dynamic):

- If a particular level of the measured confounder history through k is possible to observe in the study population (e.g. in HIV-CAUSAL), then it must also be possible to observe individuals with treatment status at k consistent with the specified strategy g.

Examples...

Example 1

Suppose:

1. We consider a deterministic static strategy "Never initiate ART" (e.g. our choice of g_2 in defining causal effect)
2. We assume (perhaps motivated by causal diagrams) that CD4 cell count history at k is needed to ensure exchangeability (it is a needed component of $\overline{L}_k$, the measured confounders)
3. Some individuals in HIV causal have CD4 measurements fall below 350 during the follow-up.
4. By clinical guidelines, any individual who has a measurement of CD4 fall below 350 initiates treatment (assume this guideline is reflected in the HIV-CAUSAL data)

This results in a positivity violation: we observe levels of CD4 cell count history in HIV-CAUSAL such that no one with that history has data consistent withe specified treatment strategy g

Example 2

Suppose everything else is the same but now, instead of considering "Never initiate ART" we consider the deterministic dynamic strategy:

- "Initiate ART in month k if CD4 first falls below 350; otherwise do not initiate at k"

In this case, the fact that, in the real world, all individuals who have a CD4 drop below 350 start treatment DOES NOT result in a positivity violation – it does not result in no one with that level of the confounder history (CD4 dropped below 350) having data consistent with the strategy. This is because we changed the strategy (i.e. the causal question).

Solution to positivity violations

So a solution to positivity violations is to change the question to one better supported by data – questions about effects of more realistic/real-world strategies

- i.e. dynamic strategies: real-world strategies are dynamic

Note: positivity is often mistakenly understood as the assumption that "there are treated and untreated individuals at each time within all levels of confounder history". This is only the definition for " always" versus 'never" treat type questions.

- Unlike exchangeability, positivity is (in principle) a testable assumption (it is an assumption about the observed data only).

Consistency Assumption

If any individual's observed treatment history is consistent with the intervention g then his observed outcome equals his counterfactual outcome Y^g

- Consistency allows us to link the observed data to counterfactual outcomes
- Sounds like a definition but it's an assumption.

Whether it holds or not depends on how we precisely define the counterfactual outcomes.

Violation of consistency

Consider simple case where we define g as "initiated treatment at baseline"

- Here Y_{60}^{g} is an individual's mortality status had, possibly contrary to fact, he initiated ART at baseline
- What if the value of this outcome depends on treatment dose and people received different doses in our study?
- For any individual in our study with $A_0 = 1$, is it guaranteed that his $Y = Y^g$?
- No if by "initiate" we mean initiate some specific dose (e.g. 20 mg)
- Yes if by "initiate" we mean "initiate according to how treatment was initiated in this study"

In some cases counterfactuals will be hard to define precisely (ill-defined); e.g. interventions on BMI...or death from competing events? Makes violations of other assumptions more likely.

The g-formula

Robins (1986,1987) showed that, given exchangeability, positivity and consistency, we can identify the outcome mean under a deterministic strategy g (static or dynamic) by the *g-computation algorithm formula* characterized by that strategy

- In recent years more simply called the *g-formula*

Like any population parameter (e.g. a population mean), the g-formula characterized by g (and contrasts in this function under different choices of g) can be estimated using some statistical procedure in a sample from the target population.

- In our example, this sample is those individuals who meet criteria for population at baseline in HIV-CAUSAL.

The (next) challenge

The g-formula characterized by some g is typically a high-dimensional function

1. driven by the dimension of the number of follow-up times $K + 1$ and
2. dimension of the measured confounder history $\overline{L}_k$, $k = 0, \ldots, K$

The g-formula for risk of a failure event by $K+1$ characterized by g

$$\sum_{\bar{l}_K} \sum_{k=0}^{K} \Pr[Y_{k+1} = 1 | \overline{L}_k = \bar{l}_k, \overline{A}_k = \overline{a}_k^g, \overline{Y}_k = \overline{C}_k = 0] \times$$
$$\prod_{j=0}^{k} \{ \Pr[Y_j = 0 | \overline{L}_{j-1} = \bar{l}_{j-1}, \overline{A}_{j-1} = \overline{a}_{j-1}^g, \overline{Y}_{j-1} = \overline{C}_{j-1} = 0] \times$$
$$f(l_j | \bar{l}_{j-1}, \overline{a}_{j-1}^g, \overline{Y}_j = \overline{C}_j = 0) \} \quad (1)$$

The g-formula for risk of a failure event by $K+1$ characterized by g

- Sum/integral over all possible levels of the measured confounder history (e.g. CD4 cell count, viral load, clinic visit history)
- The function that is summed is over is a function of time-varying conditional hazards
- Each term is conditioned on (i) a level of the confounder history and (ii) the treatment history allowed under g (for that confounder history in case of dynamic g)
- The sum over all confounder histories is weighted by a function of the joint distribution of the time-varying confounders
  - specifically, a product over all times k of the chance of seeing that level of the confounders at that time in the population conditional on that level of confounder history and treatment history consistent with g through previous time

In other words...

It's a messy function.

- So even if we could reasonably argue our "causal assumptions" (exchangeability, positivity, consistency) hold, additional constraints are needed in practice to reduce the dimension of the estimation problem (parametric models or machine learning)
- Different estimators impose constraints on different features of the data

Inverse probability weighting (IPW), parametric g-formula and TMLE are three such estimators (and more!) and within each, there are different implementations.

- Choice of estimator and implementation usually involves a bias/variance tradeoff.
  - more reliance on data versus more reliance on model constraints

Inverse probability weighting

- IPW estimators are *semiparametric* estimators: they do not require parametric specification of the whole observed data likelihood
- Any implementation of IPW requires consistent estimation of the *observed treatment mechanism*
  - Survival outcomes: density/probability function for treatment at each time in the observed population conditional on past measured treatment/confounders, surviving and remaining uncensored
  - Our example: chance of any level of treatment initiation status by each k ($A_k = 1$ or $A_k = 0$) conditional on this past
  - Denote: $f^{obs}(A_k|\overline{L}_k, \overline{A}_{k-1}, Y_k = C_k = 0)$
  - Special term: When we plug in $A_k = 1$, we call $f^{obs}(A_k|\overline{L}_k, \overline{A}_{k-1}, Y_k = C_k = 0)$ the *propensity score* at k
- Also requires consistent estimation of the *censoring mechanism* in censored data

Motivation for IPW

The reason that IPW estimators "work" (i.e. may be consistent/converge to the true value of the g-formula) is that

- The g-formula characterized by an intervention g has an equivalent IPW representation

IPW representation of the g-formula

g-formula characterized by g equivalently written as a function of weighted discrete hazards

$$\sum_{k=0}^{K} \lambda^g(k) \prod_{j=0}^{k-1} [1 - \lambda^g(j)]$$

where weighted hazards at each time are defined as

$$\lambda^g(k) = \frac{E[Y_{k+1}(1 - Y_k)W_k(g)]}{E[(1 - Y_k)W_k(g)]}$$

with "E" meaning "expectation" (population average) and weights $W_k(g)$ specific to the choice of g and the time k. Weights also depend on the person's values of treatment, confounder and censoring history....

Weights $W_k(g)$

For each person in the population who has not yet failed by k, $W_k(g)$ is defined at time k and for particular strategy of interest g as a ratio:

- numerator: Takes the value 0 if the individual has a treatment history inconsistent with g from baseline through any follow-up time k *or* if that individual is censored in that period; otherwise takes value 1
  - E.g. if g is "initiate ART at k if CD4 first drops below 450; otherwise do not initiate at k" then any individual in our HIV-infected target population who, at time k, has never had a CD4 value drop below 450 but starts treatment at k gets a 0 in time k weight numerator.
- denominator: Is a product over all times $j = 0, \ldots, k$ of $f^{obs}(A_j|\overline{L}_j, \overline{A}_{j-1}, Y_j = C_j = 0) \times \Pr[C_{j+1} = 0|\overline{L}_j, \overline{A}_{j-1}, Y_j = C_j = 0)$

A simple translation of W_k^g

All of this complicated mess can be boiled down to the following:

- For a person doing the "wrong thing" anywhere between baseline and time k according to the treatment rule g:
  - W_k^g is zero
- For a person doing the "right thing" at all times between baseline and time k according to the treatment rule g:
  - W_k^g is 1 divided by that chance of "doing the right thing" (having data consistent with g at k) among people in the observed population with that person's measured past.

IPW estimators

This result (existence of IPW representation of g-formula) gives a fairly simple approach to estimating contrasts in g-formulas characterized by a g_1 versus a g_2 through five basic steps

1. Creation of person-time data set
2. Make copies for each strategy g considered
3. Artificial censoring as soon as data inconsistent with g
4. computation of inverse probability of censoring weights
5. computation of weighted hazards and cumulative risks

Start with reviewing basic implementation of this approach and then some extensions/modifications/special cases

Example

Suppose interest is in estimating effect of dynamic strategies

- g_1: If CD4 cell count first drops below 450 at time k then initiate treatment at k; otherwise do not initiate at time k versus
- g_2: If CD4 cell count first drops below 350 at time k then initiate treatment at k; otherwise do not initiate at time k

Basic IPW: Step 1

Step 1: Construct a person-time data set

- Columns: *id*, k (follow-up month), A_k, L_k (one column per component – CD4, viral load, etc.), C_{k+1}, Y_{k+1}
- First line on which $C_{k+1} = 1$ is last record in the data for that *id* (and Y_{k+1} on that line coded as missing)
- First line on which $Y_{k+1} = 1$ is last record in the data for that *id*
- Any individual who is not censored by $K + 1$ and does not fail has $K + 1$ records with Y_{k+1} and C_{k+1} columns all 0
- In addition to these columns, will generally have additional columns with functions of "history" treatment and confounders
- *id* takes values $1, \ldots, n$ with n number meeting eligiblity for the population at baseline
- time-fixed baseline confounders are repeated on each line

This step does not depend on how we defined g_1 or g_2 (i.e. the causal question)

Illustration of subset of columns for one id

	id	k	A	CD4	AIDS	C	Y
1:	1	0	0	459	0	0	0
2:	1	1	0	552	0	0	0
3:	1	2	0	456	0	0	0
4:	1	3	1	430	0	0	0
5:	1	4	1	550	0	0	0
6:	1	5	1	552	0	0	0
7:	1	6	1	499	0	0	1

Basic IPW: Step 2

Step 2: Make copies

- Make 2 copies of this person-time data set
- Copy 1 indexed by g_1
- Copy 2 indexed by g_2

Basic IPW: Step 3

Step 3: Artificial censoring

- In each g-specific copy, create an indicator of artificial censoring
- An individual will be artificially censored in the first interval his data becomes inconsistent with g
- In this copy, at time he is artificially censored, this is his last record in the copy (even if had more records in original data)
- So same id may have different number of records, in original data, copy g_1 and copy g_2
- E.g. for the g_2 copy (initiate as soon as CD4 drops below 350; otherwise don't initiate), an individual who initiates at a time k when CD4 has not first fallen below 350 is censored at that time.
- Similarly, an individual who does not initiate at a k when CD4 has first fallen below 350 is censored at that time.

Illustration of artificial censoring in copy for g_2

	id	k	A	CD4	AIDS	C	Y
1:	1	0	0	459	0	0	0
2:	1	1	0	552	0	0	0
3:	1	2	0	456	0	0	0
4:	1	3	1	430	0	0	0
5:	1	4	1	550	0	0	0
6:	1	5	1	552	0	0	0
7:	1	6	1	499	0	0	1

Artificially censored at $k = 3$ because started before CD4 fell below 350

	id	k	A	CD4	AIDS	C	Y	Cg350
1:	1	0	0	459	0	0	0	0
2:	1	1	0	552	0	0	0	0
3:	1	2	0	456	0	0	0	0
4:	1	3	1	430	0	0	0	1

Same id in copy g_1

	id	k	A	CD4	AIDS	C	Y	Cg450	
1:	1	0	0	459	0	0	0		0
2:	1	1	0	552	0	0	0		0
3:	1	2	0	456	0	0	0		0
4:	1	3	1	430	0	0	0		0
5:	1	4	1	550	0	0	0		0
6:	1	5	1	552	0	0	0		0
7:	1	6	1	499	0	0	1		0

Same individual is not artificially censored in the copy g_1; he started treatment when his CD4 first fell below 450 which is consistent with g_1

Basic IPW: Step 4

Step 4: construct inverse probability of censoring weights.

- This is most involved step

Weight construction

In each copy, estimate the probability of NOT being artificially censored at each time (doing the right thing) conditional on (i) past treatment and confounder history and (ii) previous survival and previously free of all censoring (artificial and "real")

- Don't ever have to worry about (ii) by the data construction – because people don't have records past censoring/failure, all estimates are inherently conditioned on (ii)
- For (i) past treatment and confounder history means confounder history through k and treatment through $k-1$ $(\overline{L}_k, \overline{A}_{k-1})$ because, given this past, artificial censoring status is fully determined by your current treatment status at k
  - that is, we are estimating the probability that treatment A_k takes the right value under g given past measured confounders and previous adherence to g

Weight construction

Because the measured confounder history $\overline{L}_k$ is high dimensional (e.g. CD4 is a continuous variable) we need to make additional assumptions to estimate this probability. One simple approach:

- Pooled logistic regression
- Dependent variable indicator of artificial censoring (e.g. Cg450)
- Independent variables: function of k, current values of confounders on each line k, and past values of confounders through $k-1$ (need to create more columns for this – e.g. variable with cumavg of CD4 through each time, lagged values)

Incorporating deterministic knowledge of this probability

In our example, if you initiated by the previous time $k-1$ and were not artificially censored (you are still in the data) then you can't be artificially censored at a future time so I don't need to include you in the model (I know your probability of being uncensored at all future times is 1).

- Logistic regression can be restricted to person-time records with $A_{k-1} = 0$

Sample R code

```
modelg450<-
glm(Cg450~k+k^2+k^3+CD4+CD4lag1+AIDS+cumavgrna +blage,
data=copy450restrict,
                        family=binomial(link = "logit"))
```

Assume data set "copy450restrict" is restricted to those with treatment on previous line (lagged value of A) equal to 0. Pooled over time model assumed cubic function of time, linear function of cumavg of rna and baseline age, linear CD4 and lagged value of CD4. These are restrictive assumptions:

- Could alternatively use Superlearner (see Maya's workshop) to estimate these probabilities. E.g. can include different candidate functional forms in the library.

Artificial censoring weight construction for each record in each copy

Next, estimate the probability of NOT being artificial censored in each data copy at each time k for each id *who is uncensored on line* k. In our example

- For those with previous value of initiation indicator 1 (they started by an earlier time, $A_{k-1} = 1$), set this probability to 1 (by previous arguments)
- For those with previous value of initiation indicator 0 (they didn't start by previous month, $A_{k-1} = 0$), predict this probability from the model using his observed covariate history (e.g. predict in R)

Call each of these estimated probabilities for each person-time $\hat{p}^g_{i,k}$

Artificial censoring weight construction for each record in each copy

- For each record that is not artificially censored on line k, the artificial censoring weight is defined as the cumulative product of inverse probabilities $1/\hat{p}^g_{i,0} \times ...1/\hat{p}^g_{i,k}$ (on first line this is just $1/\hat{p}^g_{i,0}$)
- For each record that is artificially censored on line k, this weight is set to zero.

Actual censoring weights

When there is actual censoring in the original data, we additionally construct weights for this type of censoring. Analogous process

- Estimate probability of actual censoring at each time given treatment and confounder history (all through k), previously surviving and remaining uncensored
- E.g. pooled over time logistic regression with dependent variable actual censoring indicator
- Independent variables function of treatment and confounder history

Actual censoring weight construction for each record in each copy

Next, estimate the probability of NOT being actually censored in each data copy at each time k for each id *who is uncensored on line* k. In our example

- Predict from model for these records (predict in R)

Call each estimated probability for each person-time record $\hat{p}^c_{i,k}$

- The actual censoring weight for these records is defined as the cumulative product of inverse probabilities $1/\hat{p}^c_{i,0} \times ...1/\hat{p}^c_{i,k}$ (on first line is just $1/\hat{p}^c_{i,0}$)

The actual censoring weight for records who *are artificially censored on line* k *is set to zero*.

Final weight for each id and time k in copy g

Take product of actual and artificial censoring weights for each id on line k

$$1/\hat{p}^{g}_{i,0} \times ...1/\hat{p}^{g}_{i,k} \times 1/\hat{p}^{c}_{i,0} \times ...1/\hat{p}^{c}_{i,k}$$

Steps 1 to 4 recap

1. Construct person-time format data (not specific to the question)
2. Create two copies of original person-time data (one indexing each hypothetical strategy)
3. Artificially censor each copy according to that copy specific rule (there also may be actually censoring of people who are not artificially censored)
4. Construct person-time weights for each copy and each person-time record in each copy
   - records that are either artificially or actually censored at k get a weight of zero on line k
   - those not, get a weight that is an inverse probability of censoring weight based on estimates of censoring probabilities (this will be another time-varying variable in each artificially censored data set)

Basic IPW: Step 5

In each copy, implement complement of weighted Kaplan-Meier as function of weighted time-varying hazards

$$\sum_{k=0}^{K} \hat{\lambda}^g(k) \prod_{j=0}^{k-1} [1 - \hat{\lambda}^g(j)]$$

where each $\hat{\lambda}_k^g$ is just a fraction: denominator is sum of the weights of all time k records in that artificially censored copy g and numerator is sum of the weights in the subset of those whose failure (death) indicator is 1.

- Repeat in each copy and take difference/ratio for effect estimate
- For 95% confidence intervals, repeat Steps 1 through 5 in many (e.g. 1000) resamples with replacement of the original n ids, sort effect estimates and take 2.5 and 97.5 percentiles as lower and upper bounds.

Stabilized weights

The IP of censoring weights in this basic implementation are simply estimates of the weights W_k^g that we defined earlier in defining the IPW representation of the g-formula.

- They can be highly variable, and estimator can perform poorly, particularly if there are *near positivity violations*
  - Certain individuals are observed to follow g in the data, but they are rare (few people in the data with their treatment and covariate history adhere to g).
- In this case, the weights we described are 1 over a tiny value

Stabilized weights can help.

Stabilized weights

Stabilized weights effectively multiply weights W_k^g by a "constant" function.

- One convenient constant that doesn't substantially change the basic algorithm is the product from $j = 0, \ldots, k$ of the probabilities of being uncensored at each time j given previous survival and remaining uncensored (constant because it is a feature of the population and not specific to individual's confounder and treatment values)
- can estimate these probabilities by proportions in the sample data
- May not stabilize "enough" (can look at weight distribution, for example, is it skewed?) – also how off is estimate from that obtained using other methods (more on this later)

Stabilized weights and marginal structural models

- Can potentially further reduce extreme weights by changing more features of the algorithm – but at the expense of more constraints on the observed data distribution.
- For example, a modification of the algorithm changes Step 5 so that, instead of weighted Kaplan-Meier we impose a *marginal structural model* (MSM) on the λ_k^g functions conditional on some subset of the baseline covariates L_0 (call this V), estimate the coefficients of that model and then transform back to risk.
- E.g. for V selected as race and sex we assume

$$\text{logit}\{\lambda^g(k, V)\} = \beta_0 + \beta_1 k + \beta_2 I(g = g_1) + \beta_3 race + \beta_4 sex$$

Stabilized weights and marginal structural models

- Stack the two artificially censored data copies indexed by g_1 and g_2 creating one long data set. Add a new column to this new stacked data set g which takes value 0 if records come from copy g_2 and 1 if records come from copy g_1.
- Fit a weighted pooled over time logistic regression outcome model with: dependent variable the event indicator Y_{k+1} and independent variables a function of the copy indicator g, time index k and V (could choose $V = L_0$)
- Weights are defined as in basic algorithm but with each term in cumulative product multiplied by estimated censoring probabilities conditional only on V (as opposed to the censoring probabilities in the denominator which condition on time-varying covariates $\overline{L}_k, \overline{A}_{k-1}$).

Stabilized weights and marginal structural models

This procedure targets the weighted hazard functions λ_k^g but now conditional (within levels of) V, functions of this subset are constants (this is why we can multiply numerator by this function even though V varies in the population). These weights tend to get less variable as more components of L_0 included in V.

- Special case where we choose V empty, we are back to the previous stabilization approach and don't need to make an MSM assumption (but weights likely more variable)

Stabilized weights and marginal structural models

Can still get back population-level (marginal over V) effect estimates on risk scale by computing conditional risks from the hazards and averaging out V:

- From the estimated coefficients of this weighted outcome model, predict for each of the n individuals at baseline the hazard at k under g for their values of V – $\hat{\lambda}^g(k, V_i)$
- From these model based hazard estimates apply the same kaplan meier computation to estimate the risk by $K+1$ for each individual (each V_i in the data) under each g and average these individual g-specific risk estimates to get the population risk estimate under each g.

$$\frac{1}{n}\sum_{i=1}^{n}\sum_{k=0}^{K}\hat{\lambda}^g(k, V_i)\prod_{j=0}^{k-1}[1-\hat{\lambda}^g(j, V_i)]$$

For effect estimates take difference/ratio under each g. Can bootstrap for CIs. .

Marginal structural models

- In any of the previously described implementations (with or without stabilized weights), by considering only two choices of g, a lot of people might be artificially censored.
- So even if n is large, very few people may actually contribute to the weighted estimator (a lot of zero weights).
- Can impose even stronger MSM assumptions to use more data at the expense of more potential for model misspecification bias

Marginal structural models

Consider dynamic strategies of the form

- If CD4 cell count first drops below x at time k then initiate treatment at k; otherwise do not initiate at time k

We initially considered only two choices of x (350 or 450). What if we considered many choices (e.g. $x = 450, \ldots, 200$ in increments of 50). Denote these $g(x)$.

Marginal structural models

Can use the copy and stacking approach we just considered but now

- We make more than 2 copies (there is a copy for each level of x considered)
- Artificial censoring and weight creation for each $g(x)$ as before
- Weighted outcome regression to estimate hazards at each time k under each $g(x)$ (maybe conditional on V for better stabilized weights) can be a function of k and x (and possibly V).
- The more levels of x (the more copies), the stronger assumptions this model is making; e.g. if we just put in x then weighted logistic outcome regression then we are assuming MSM is linear in x
- If assumption is right, we can increase precision for $g(x) = 350$ versus $g(x) = 450$ but if wrong increase bias

But this model answers more questions (if it is correct). Could report risk curves as a function of x.

Marginal structural models and static deterministic strategies

When we restrict attention to static interventions g, a very natural MSM arises. In our example, all such interventions are covered by:

- Do not initiate ART
- Initiate ART at time k, for any level of k

Every individual has data consistent with one of these interventions – the one corresponding to their observed treatment pattern $\overline{A}_k$.

MSMs and static deterministic strategies

So everyone's follow-up gets used in this case (no one is artificially censored for all of these interventions, they will always be following at least one). An MSM that smooths over all of these can be fit as follows:

- fit pooled logistic with dependent variable Y_{k+1} and independent variables function of $\overline{A}_k$ (observed treatment history through k) and k itself (also possibly subset of V for stabilized weights)
- Remaining steps to get risks under each g can be implemented as previously described

Potentially strong model assumption. And causal questions which are contrasts in risks under static strategies don't correspond to real-world strategies which are dynamic (positivity/near positivity violations more likely). See Hernán et al. (2000)

Grace periods

- Another way to improve precision/use more data (with or without an MSM) is to allow a "grace period"
- For example could consider strategies
  "If CD4 cell count first drops below x at time k then initiate treatment within m months of k; otherwise do not initiate at time k". E.g. $m = 6$
- Could choose two choices of x or impose marginal structural models, etc. within this new question.
- For any implementation, fewer individuals with be artificially censored as m gets larger because it is a more flexible strategy.
- Same broad approach as without grace period: no one is artificially censored during grace period because intervention says you can do what you want in that period – so probability of being (artificially uncensored in this period) is 1 (incorporate into weight construction as above).

Dual Interpretation

These interventions with grace periods can be understood in two ways:

- During grace period, do not intervene
- During grace period, define the treatment rule as "assign treatment initiation according to a random draw from the observed treatment density $f^{obs}(A_k|\overline{L}_k, A_{k-1} = 0, C_k = Y_k = 0)$

Latter is an example of a random dynamic strategy. Second requires stronger exchangeability assumption but may be more desirable for generalizing results to future populations.

IPW and random dynamic strategies in general

- Artificial censoring approach does not work for IPW estimation of effects of random dynamic strategies in general.
- In general similar approach but weight denominator is estimate of observed treatment density, weight numerator an intervention density (the chance of receiving a level of treatment under the hypothetical strategy)
- Can show that weights constructed under artificial censoring approach for deterministic interventions/grace periods are estimating this intervention to observed data density ratio (when weight denominator models are correct)
- Because the chance of receiving any level of treatment in the observed data is probabilistic not deterministic, the weights used in IPW for random dynamic strategies can inherently be more stable
- E.g. can define the intervention density to be just a small shift of the observed density (e.g. see Munoz-Diaz and van der Laan,

Assumptions for random dynamic strategies

We only reviewed exchangeability, positivity and consistency (identifying) assumptions for deterministic strategies g

- Can show that risk of event of interest by $K+1$ (e.g. 5-year mortality risk) under a hypothetical random dynamic strategies is equal to a weighted average of risks under a subset of deterministic treatment strategies
- Subset includes any deterministic strategy that could be observed if we implemented that random dynamic strategy
- Formally, subset includes any strategy g under which positivity would hold if the the observed treatment density were replaced with the intervention density for that hypothetical strategy.

Therefore, the risk under a hypothetical random dynamic strategy is identified in a study under which exchangeability, positivity and consistency hold for all deterministic g in that subset (see for example Young et al. (2014) in reference list).

Parametric g-formula

Parametric g-formula is another approach that can be used to estimate effects of time-varying treatment strategies:

- Requires consistent estimation of full observed data likelihood
- It provides a sparse data "solution" in face of near positivity violations/few adherent to strategies over time but, again, at the expense of potential model missspecification
- On the whole, model assumptions are arguably even more extreme than those imposed by the MSM
- g-null paradox – under no causal treatment effect at any time but t-v confounding affected by past treatment, some model misspecification is guaranteed under standard parametric models (except in unlikely cases)

Parametric g-formula

Why do we use it?

- Up until a few years ago, we did not know how to construct semi-parametric estimators (e.g. IPW, TMLE) for certain types of dynamic strategies on continuously measured time-varying treatments
  - Interventions that maintain a continuous treatment with a range via interventions that depend on the *natural value of treatment*–Richardson and Robins (2013)
- Young et al. (2014) first showed how an IPW estimator could be derived for these interventions
- Currently still no extension for these particular interventions to TMLE

Parametric g-formula

Why do we use it?

- Implementing two methods that rely on constraints on different components of the observed data distribution can be confirmatory of our conclusions
- Drastically different results may suggest poor data support for the question(near positivity violations, few individuals following strategies) – this is informative in itself

TMLE is clearly becoming another option with many recent developments in practical implementation (e.g. stremr package, 'long-format TMLE"; Sofrygin, O., van der Laan, M. J., and Neugebauer, R. (2016). stremr: Streamlined Estimation of Survival for Static, Dynamic and Stochastic Treatment and Monitoring Regimes. R package version 0.31.)

Parametric g-formula

- Another potential advantage of parametric g-formula? Allows incorporation of a priori knowledge of distribution of time-varying confounders (see Young et al., 2011). Even when such knowledge is available may not be enough to outweigh the areas of the data where no knowledge is available.
- "Computational advantage": algorithm changes little for different strategies (unlike IPW which can be "very strategy specific" in implementation)

Parametric g-formula algorithm

Estimates the g-formula for risk by $K+1$ under g by

1. Estimate each component of the g-formula using regression models (under distributional assumptions)
   - hazards at each time given past treatment and confounders
   - joint distribution of the confounders at each time given past treatment and confounders
2. Under confounders models and starting with each of the n values of L_0, simulate confounders at each time k and assign treatment under g
3. Under the hazard model estimate the hazard at each k for each of these n simulated "histories" under g
4. Using complement of Kaplan-Meier, estimate for each of these n histories the risk by $K+1$ under g
5. Population risk under g – average of these n history specific risks

Repeat for each g and take difference/ratio. Bootstrap for 95% CIs.

Software and other materials

- GFORMULA SAS macro with documentation and examples can be accessed here: `https://github.com/CausalInference/GFORMULA-SAS`
- CIMPOD 2017 workshop materials with worked SAS example for CD4-based dynamic strategies in HIV-CAUSAL (slides and code) can be accessed here: `https://drive.google.com/file/d/0B_BEunAWYES5Si1VUENXSVNCSjg/view`
- Youtube video of 2017 workshop: `https://www.youtube.com/watch?v=evFBLp4MekI` (sound in beginning is bad but gets better)
- R package...coming very soon (technical report with many worked examples will be made available on arXiv.org)

Selected References

Cain, L. E., Robins, J. M., Lanoy, E., Logan, R., Costagliola, D., and Hernán, M. A. (2010). When to start treatment? a systematic approach to the comparison of dynamic regimes using observational data. *International Journal of Biostatistics*, 6. Article 18.

Hernán, M. A., Brumback, B., and Robins, J. M. (2000). Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology*, 11(5):561–570.

Hernán, M. A., Lanoy, E., Costagliola, D., and Robins, J. M. (2006). Comparison of dynamic treatment regimes via inverse probability weighting. *Basic & Clinical Pharmacology & Toxicology*, 98:237–242.

Li, X., Young, J., and Toh, S. (2017). Estimating Effects of Dynamic Treatment Strategies in Pharmacoepidemiologic Studies with Time-varying Confounding: A Primer. *Current Epidemiology Reports*, 4(4):288–297.

Richardson, T. S. and Robins, J. M. (2013). Single World Intervention Graphs (SWIGs): a unification of the counterfactual and graphical approaches to causality. *Center for Statistics and the Social Sciences, University of Washington Series*. Working Paper Number 128.

Robins, J. M. (1986). A new approach to causal inference in mortality studies with a sustained exposure period: application to the healthy worker survivor effect. *Mathematical Modelling*, 7:1393–1512. [Errata (1987) in Computers and Mathematics with Applications 14, 917-921. Addendum (1987) in Computers and Mathematics with Applications 14, 923-945. Errata (1987) to addendum in Computers and Mathematics with Applications 18, 477.].

Robins, J. M. and Wasserman, L. (1997). Estimation of effects of sequential treatments by reparameterizing directed acyclic graphs. In Geiger, D. and Shenoy, P., editors, *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, pages 409–420. San Francisco: Morgan Kaufmann.

Young, J. G., Cain, L. E., Robins, J. M., O'Reilly, E. J., and Hernán, M. A. (2011). Comparative effectiveness of dynamic treatment regimes: an application of the parametric g-formula. *Statistics in Biosciences*, 3(1):119–143.

Young, J. G., Hernán, M. A., and Robins, J. M. (2014). Identification, estimation and approximation of risk under interventions that depend on the natural value of treatment using observational data. *Epidemiologic Methods*, 3(1):1–19.

Zhang, Y., Young, J., Thamer, M., and Hernán, M. (2018). Comparing the Effectiveness of Dynamic Treatment Strategies Using Electronic Health Records: An Application of the Parametric g-Formula to Anemia Management Strategies. *Health Services Research*, 53(3):1900–1918.

Types of paths on a Causal DAG

In understanding how to read exchangeability off of a causal DAG, useful to understand different types of paths on a Causal DAG:

1. Causal path
2. Unblocked backdoor path
3. Paths containing common effects (colliders) of two variables along the path

These three types of paths represent different underlying causal structures that give rise to association between variables in your data.

Causal Path

$$A_0 \longrightarrow Y$$

Path $A_0 \rightarrow Y$ is called a "causal path". E.g. allows that baseline initiation causes survival status by 60 months. If this is a causal DAG assume no common causes of treatment and outcome (would be true in a trial with baseline randomization and perfect baseline adherence).

Unblocked backdoor path

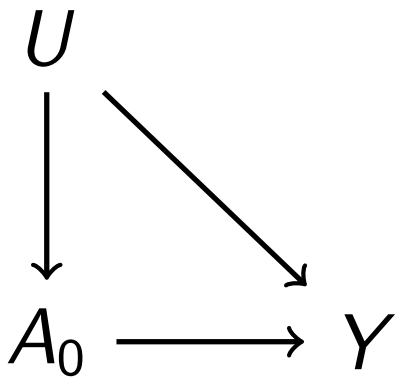

Path $A_0 \leftarrow U \rightarrow Y$ "unblocked backdoor path". Allows that baseline initiation status and later mortality share a common cause.

- Conditioning on a variable along an unblocked backdoor path, blocks the path.

Here we would say "Conditional on U there are no unblocked backdoor paths between A and Y"

Paths containing colliders

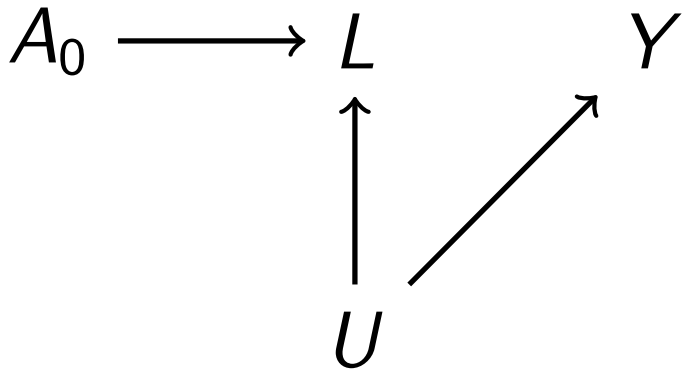

L is a common effect of A_0 and U. Say L is a "collider" on the path $A \rightarrow L \leftarrow U \rightarrow Y$ (L could be interim CD4 cell count). Paths containing colliders are blocked.

- This graph could again represent a trial with random assignment of A

Conditioning on a collider

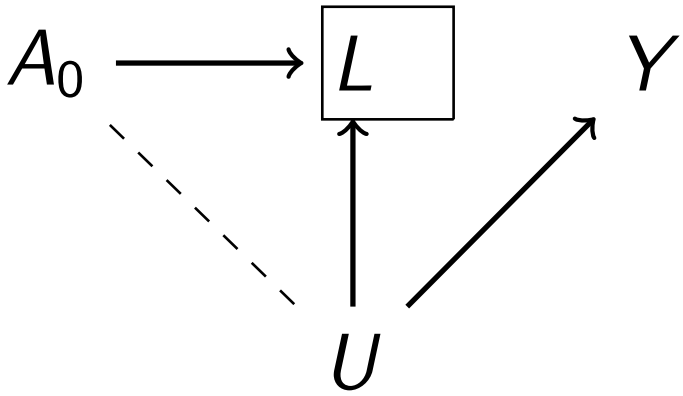

Conditioning on a collider on a path between A_0 and Y opens that path and associates A_0 and Y (through U). Why?

- Knowing patient has low interim CD4 count and in initiator group, makes them more likely to have had the virus longer at baseline (they started out sicker).

That is, knowing the effect, makes it more likely that, if you didn't have one cause, you had the other.

Graphical evaluation of exchangeability

Exchangeability is specific to both the intervention g and time k. To evaluate exchangeability $Y^g \coprod A_k | \overline{L}_k, \overline{A}_{k-1} = \overline{a}^g_{k-1}$ for a given strategy g and time k on a DAG G, manipulate G as follows:

1. Remove all arrows out of A_k
2. Remove all arrows into $A_{k+1}, \ldots, A_K$
3. For $j = k+1, ..., K$ add back arrows from any component of L_s, $s \leq j$ into A_j if A_j is assigned based on that component of L_s under strategy g
   - English translation: if under the strategy g a measured covariate value would be used to determine treatment assignment, add an arrow from that covariate to treatment – e.g. under a dynamic strategy where CD4 history (a component of L_k) is used to determine whether treatment should be initiated.

Call this manipulated graph G^g_k.

Graphical evaluation of exchangeability

If the DAG G is a causal DAG then $Y^g \coprod A_k | \overline{L}_k, \overline{A}_{k-1} = \overline{a}^g_{k-1}$ holds if

- there are *no unblocked backdoor paths* between A_k and Y conditional on $\overline{L}_k, \overline{A}_{k-1}$ on G^g_k.

Example for $K = 1$

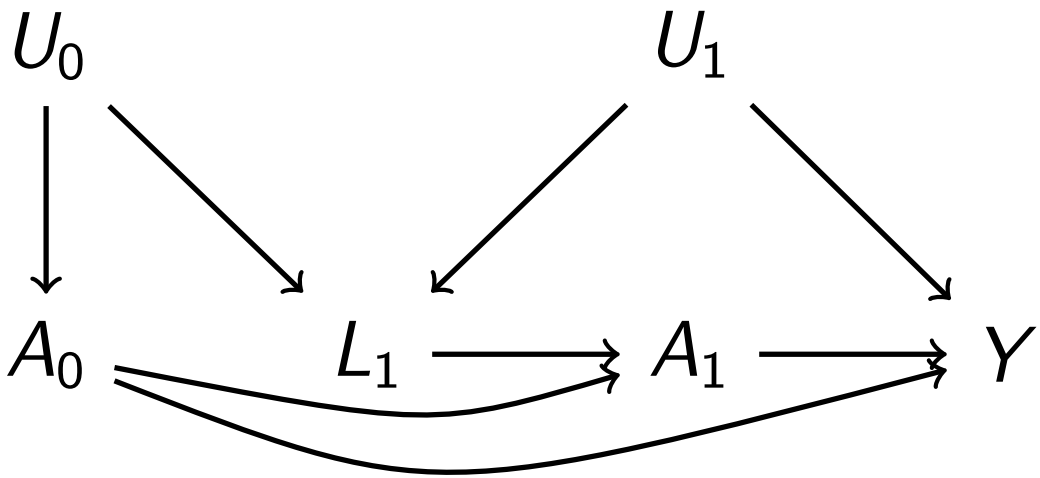

Evaluate $Y^g \coprod A_k | \overline{L}_k, \overline{A}_{k-1} = \overline{a}^g_{k-1}$ for $k = 0, 1$, static $g = (a_0, a_1)$ and dynamic $g = (a_0, a_1^g = l_1)$. Modified from Robins and Hernan chapter "Estimation of the causal effects of time-varying exposures"

Construct G_1^g for static g

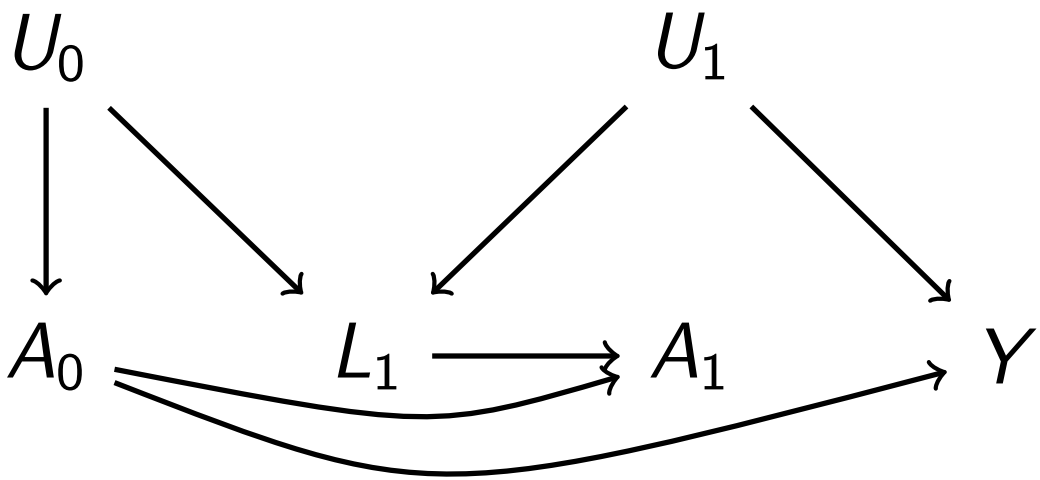

Construct G_1^g for dynamic g

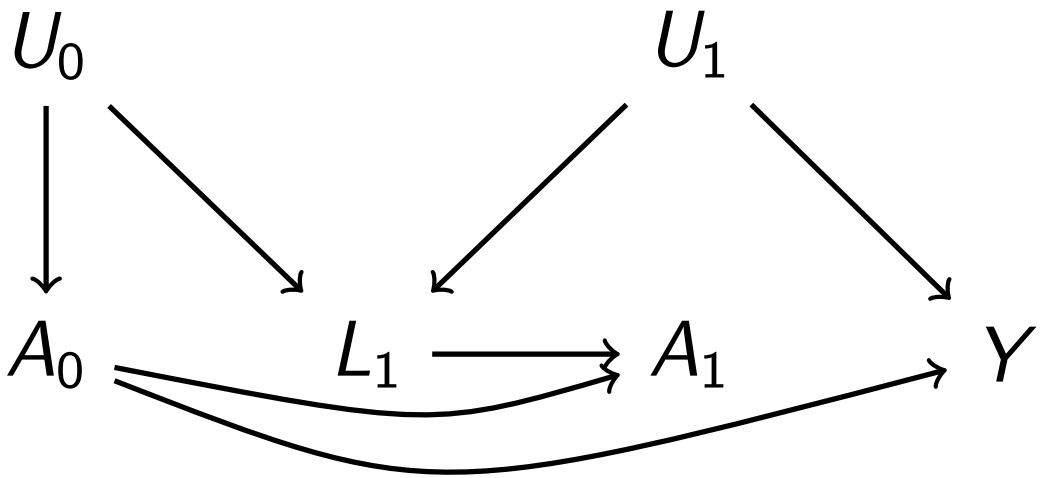

same

Construct G_0^g for static g

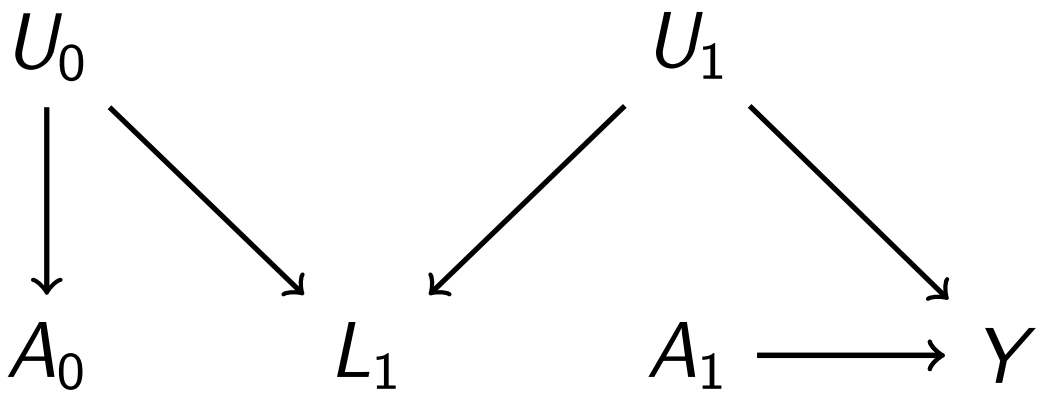

Construct G_0^g for dynamic $g = (a_0, a_1^g = l_1)$