

Modern Statistical Learning Methods for Observational Biomedical Data

Chapter 5: Identification and inference on the average treatment effect of a time-varying intervention

David Benkeser
Emory University

Marco Carone
Univ. of Washington

MODULE 8

Summer Institute in Statistics for Clinical and Epidemiological Research

July 2023

Contents of this chapter


- 1 Motivation for considering time-varying interventions
- 2 Identification of average treatment effects
- 3 Failure of naive approach to causal inference
- 4 Simple estimation procedures
- 5 Improved estimation procedures
- 6 Illustrations in practice
- 7 Dynamic treatment rules

Motivation for considering time-varying interventions

In many clinical contexts, the treatment of interest is administered in phases over time.

- antihypertensive drug therapy administered daily;
- biphosphonate drug therapy administered weekly;
- injection of antiretroviral suspension administered every month;
- immunosuppressant infusion therapy administered every two months.

The observed data is often of the form



where we have defined components

- L_k = covariates recorded at time k ;
- A_k = treatment assignment at time k ;
- Y = outcome recorded at the end of the study.

Motivation for considering time-varying interventions

We can consider the counterfactual outcome $Y(a_0, a_1, \dots, a_K)$ defined by enforcing treatment assignment $(A_0, A_1, \dots, A_K) = (a_0, a_1, \dots, a_K)$.

This allows to define causal contrasts that address the scientific question of interest.

(Chapters 24-26 of van der Laan & Rose, 2011; Chapter 4 of van der Laan & Rose, 2018; Chapter 19 of Hernán & Robins, 2018)

Weekly alendronate therapy for osteoporosis and one-year incidence of hip fracture:

- L_k = covariates recorded at week k
(e.g., sex, age, dexascan values, thyroid hormone levels, side effects, fracture status);
- A_k = indicator that alendronate was taken at week k ;
- Y = indicator that hip fracture occurred within one year.

We may be interested in the average effect

$$E[Y(1, 1, \dots, 1)] - E[Y(0, 0, \dots, 0)]$$

of year-long weekly alendronate therapy on one-year risk of hip fracture versus no alendronate therapy, or other contrasts defined by values of $(a_0, a_1, \dots, a_{52})$.

Motivation for considering time-varying interventions

Even when the treatment is administered at a single time-point, it is often the case that the data are incompletely recorded in the follow-up period.

- **missing data:** patient did not show up to a scheduled clinic visit;
- **loss to follow-up:** patient moved out of the country and dropped out of the study.

It would be natural then to consider a counterfactual outcome defined by enforcing

- 1 the administration of a particular treatment (baseline only or time-varying);
- 2 complete follow-up and complete recording of data (time-varying).

What would the outcome have been had:

- > the patient taken an experimental treatment regime, the follow-up been complete, and all data been completely recorded?
- > the patient taken a control treatment regime, the follow-up been complete, and all data been completely recorded?

Motivation for considering time-varying interventions

For example, if treatment is only administered at baseline, we could set:

L_k = covariates recorded at time k ;

A_0 = treatment assignment at time 0 (i.e., at baseline);

A_k = indicator that, at time k , patient has not yet been lost to follow-up and all measurements on this patient are complete;

Y = outcome recorded at the end of the study.

We might then be interested in

$$ATE = E[Y(1, 1, 1, \dots, 1)] - E[Y(0, 1, 1, \dots, 1)] .$$

Motivation for considering time-varying interventions

If treatment is administered over time, we could instead set:

- L_k = covariates recorded at time k ;
- $A_{k,1}$ = indicator that, at time k , patient has not yet been lost to follow-up and all measurements on this patient are complete;
- $A_{k,2}$ = indicator of treatment assignment at time k ;
- Y = outcome recorded at the end of the study.

and let $Y(\underbrace{(a_{0,1}, a_{0,2})}_{a_0}, \underbrace{(a_{1,1}, a_{1,2})}_{a_1}, \dots, \underbrace{(a_{K,1}, a_{K,2})}_{a_K})$ be the counterfactual defined by

$$(A_0, A_1, \dots, A_K) = (a_0, a_1, \dots, a_K),$$

where we write $A_k := (A_{k,1}, A_{k,2})$.

We might then be interested in

$$E[Y((1, 1), (1, 1), \dots, (1, 1))] - E[Y((1, 0), (1, 0), \dots, (1, 0))] .$$

Identification of average treatment effects

Our goal is to contrast the mean outcome under various sequences of interventions occurring over time. We focus on treatment profile $(a_0, a_1, \dots, a_K) = (1, 1, \dots, 1)$ but other profiles can be dealt with similarly.

(Sequentially) randomized trial

We can imagine conducting a trial in which, at each of these time-points, individuals are randomized to one of the possible interventions.

In this case, at each time-point, the intervention assignment is independent of the possible counterfactual outcomes.

$$Y(1, 1, \dots, 1) \perp A_0, \quad Y(1, 1, \dots, 1) \perp A_1, \quad \dots, \quad Y(1, 1, \dots, 1) \perp A_K.$$

Identification of average treatment effects

Our goal is to contrast the mean outcome under various sequences of interventions occurring over time. We focus on treatment profile $(a_0, a_1, \dots, a_K) = (1, 1, \dots, 1)$ but other profiles can be dealt with similarly.

Observational study

In an observation study, there are often factors that influence both the intervention assignment mechanisms and the counterfactual outcome distribution.

Examples of time-varying confounding:

- a patient may discontinue chemotherapy because they have ceased to respond, which may itself be a marker of disease progression;
- a patient may have ceased smoking because they developed respiratory symptoms, which may be a sign of lung cancer.

Identification of average treatment effects

Our goal is to contrast the mean outcome under various sequences of interventions occurring over time. We focus on treatment profile $(a_0, a_1, \dots, a_K) = (1, 1, \dots, 1)$ but other profiles can be dealt with similarly.

Observational study

The vector of time-varying covariates (L_0, L_1, \dots, L_K) can be used to deconfound the relationship between Y and (A_0, A_1, \dots, A_K) provided

$$Y(1, 1, \dots, 1) \perp A_0 \mid L_0, \quad Y(1, 1, \dots, 1) \perp A_1 \mid \bar{L}_1, A_0 = 1, \dots \\ Y(1, 1, \dots, 1) \perp A_K \mid \bar{L}_K, \bar{A}_{K-1} = 1_K,$$

where the symbol 1_j is used to denote a vector $(1, 1, \dots, 1)$ of length j .

In other words, at each time-point, intervention assignment is randomized *within each stratum* defined by recorded patient history up to that point, among patients who have received the intervention of interest so far.

This is referred to as the **sequential randomization** (or **exchangeability**) **condition**.

Identification of average treatment effects

Our goal is to infer what the mean outcome would be in the target population under the multi time-point intervention of interest.

We must be able to observe the intervention of interest for each different “type” of individual (as defined by recorded covariates) from this population:

- $P(A_0 = 1 \mid L_0 = \ell_0) > 0$ for each possible ℓ_0 ;
- $P(A_1 = 1 \mid \bar{L}_1 = \bar{\ell}_1, A_0 = 1) > 0$ for each possible $\bar{\ell}_1$;
- ...
- $P(A_K = 1 \mid \bar{L}_K = \bar{\ell}_K, \bar{A}_K = 1_K) > 0$ for each possible $\bar{\ell}_K$.

As before, this is referred to as the **positivity condition**.

Identification of average treatment effects

We first focus on a setting with three time-points ($K = 2$).

$$E[Y | A_1 = 1, L_1, A_0 = 1, L_0]$$

$$= E[Y(1, 1) | A_1 = 1, L_1, A_0 = 1, L_0] = E[Y(1, 1) | L_1, A_0 = 1, L_0]$$

= mean counterfactual outcome among patients treated at time 0, with covariate value L_1 at time 1 and L_0 at time 0

$$E[E[Y | A_1 = 1, L_1, A_0 = 1, L_0] | A_0 = 1, L_0]$$

$$= E[Y(1, 1) | A_0 = 1, L_0] = E[Y(1, 1) | L_0]$$

= mean counterfactual outcome among patients with covariate value L_0 at time 0

$$E[E[E[Y | A_1 = 1, L_1, A_0 = 1, L_0] | A_0 = 1, L_0]]$$

$$= E[Y(1, 1)]$$

= mean counterfactual outcome

Identification of average treatment effects

This can be generalized to an arbitrary number of time-points (i.e., arbitrary K).

Under the sequential randomization and positivity conditions, it can be shown that the mean counterfactual outcome value $E[Y(1, 1, \dots, 1)]$ is given by

$$E\left[E\left[E\left[\dots\left[E\left[E\left(Y \mid \bar{A}_K = 1_{K+1}, \bar{L}_K\right) \mid \bar{A}_{K-1} = 1_K, \bar{L}_{K-1}\right] \dots\right] \mid \bar{A}_1 = 1_2, \bar{L}_1\right] \mid A_0 = 1, L_0\right]\right],$$

where, for any k , we write $\bar{A}_k := (A_0, A_1, \dots, A_k)$ and $\bar{L}_k := (L_0, L_1, \dots, L_k)$.

This is the multi time-point extension of the **G-computation formula** (Robins, 1986).

Identification of average treatment effects

This can be generalized to an arbitrary number of time-points (i.e., arbitrary K).

Under the sequential randomization and positivity conditions, it can be shown that the mean counterfactual outcome value $E[Y(1, 1, \dots, 1)]$ is given by

$$E \left[E \left[E \left[\dots \left[E \left[E \left(Y \mid \bar{A}_K = 1, \bar{L}_K \right) \mid \bar{A}_{K-1} = 1, \bar{L}_{K-1} \right] \dots \right] \mid \bar{A}_1 = 1, \bar{L}_1 \right] \mid A_0 = 1, L_0 \right] \right].$$

The diagram illustrates the nested structure of the expectation operator in the equation above. It consists of five horizontal green brackets of increasing length, each centered under a specific part of the nested expectation expression:

- The innermost bracket is labeled $\bar{Q}_{K+1}(\bar{L}_K)$ and is positioned under the innermost E operator.
- The next bracket is labeled $\bar{Q}_K(\bar{L}_{K-1})$ and is positioned under the next E operator.
- The third bracket is labeled $\bar{Q}_2(\bar{L}_1)$ and is positioned under the E operator for $\bar{A}_1 = 1, \bar{L}_1$.
- The fourth bracket is labeled $\bar{Q}_1(L_0)$ and is positioned under the E operator for $A_0 = 1, L_0$.
- The outermost and longest bracket is labeled \bar{Q}_0 and is positioned under the outermost E operator.

Identification of average treatment effects

The idea of inverse probability of treatment weighting naturally suggests a simple identification formula, as in the single time-point setting.

Individuals who received the entire treatment regime of interest are not representative of the target population because of (time-varying) confounding.

What about upweighting their contribution to recover representativeness?

The generalized propensity scores are defined as

$$g_k(\bar{\ell}_k) := P(A_k = 1 \mid \bar{L}_k = \bar{\ell}_k, \bar{A}_{k-1} = \mathbf{1}_k) \text{ for } k = 0, 1, 2, \dots, K.$$

For a patient with partial history $\bar{L}_k = \bar{\ell}_k$, the composite probability of receiving treatment profile $\bar{A}_k = (1, 1, \dots, 1)$ is simply given by

$$\bar{g}_k(\bar{\ell}_k) := \prod_{j=0}^k g_j(\bar{\ell}_j) = g_0(\ell_0)g_1(\bar{\ell}_1)g_2(\bar{\ell}_2) \cdots g_k(\bar{\ell}_k).$$

Identification of average treatment effects

If $\bar{g}_K(\bar{\ell}_K)$ is small, a patient with history $\bar{\ell}_K$ and treatment profile $\mathbf{1}_{K+1}$ is an unlikely occurrence in the sampling population.

This patient needs to serve as stand-in for the many such patients not seen.

The **IPTW identification formula** is given by

$$E[Y(\mathbf{1}, \mathbf{1}, \dots, \mathbf{1})] = E\left[\left\{\frac{A_0 A_1 \dots A_K}{\bar{g}_K(\bar{L}_K)}\right\} Y\right]$$

for treatment profile $(a_0, a_1, \dots, a_K) = (\mathbf{1}, \mathbf{1}, \dots, \mathbf{1})$, and similarly for other profiles.

How does the risk of positivity violations compare to the single time-point setting?

Identification of average treatment effects

The equivalence between the IPTW and G-computation identification formulas can be established through repeated uses of the law of total expectation.

$$\begin{aligned} E \left[\left\{ \frac{A_0 A_1 \dots A_K}{g_0(L_0) g_1(\bar{L}_1) \dots g_K(\bar{L}_K)} \right\} Y \right] &= E \left[E \left[\left\{ \frac{A_0 A_1 \dots A_K}{g_0(L_0) g_1(\bar{L}_1) \dots g_K(\bar{L}_K)} \right\} Y \mid \bar{L}_K, \bar{A}_K \right] \right] \\ &= E \left[\left\{ \frac{A_0 A_1 \dots A_K}{g_0(L_0) g_1(\bar{L}_1) \dots g_K(\bar{L}_K)} \right\} \bar{Q}_{K+1}(\bar{L}_K) \right] \\ &= E \left[\left\{ \frac{A_0 A_1 \dots A_{K-1}}{g_0(L_0) g_1(\bar{L}_1) \dots g_{K-1}(\bar{L}_{K-1})} \right\} \frac{\bar{Q}_{K+1}(\bar{L}_K)}{g_K(\bar{L}_K)} E(A_K \mid \bar{L}_K, \bar{A}_{K-1}) \right] \\ &= E \left[\left\{ \frac{A_0 A_1 \dots A_{K-1}}{g_0(L_0) g_1(\bar{L}_1) \dots g_{K-1}(\bar{L}_{K-1})} \right\} E[\bar{Q}_{K+1}(\bar{L}_K) \mid \bar{L}_{K-1}, \bar{A}_{K-1} = \mathbf{1}_{K-1}] \right] \\ &= E \left[\left\{ \frac{A_0 A_1 \dots A_{K-1}}{g_0(L_0) g_1(\bar{L}_1) \dots g_{K-1}(\bar{L}_{K-1})} \right\} \bar{Q}_K(\bar{L}_{K-1}) \right] = \dots \end{aligned}$$

Failure of naive approach to causal inference

When can causal effects be read off regression models in single time-point settings?

LINEAR MODEL without an interaction between A and W :

$$E(Y | A = a, W = w) = \beta_0 + \beta_1 a + \beta_2 w$$

$$E(Y | A = 1, W = w) - E(Y | A = 0, W = w) = \beta_1$$

$$E[E(Y | A = 1, W) - E(Y | A = 0, W)] = \beta_1$$

Regression coefficients generally cannot be interpreted as average treatment effects.

Failure of naive approach to causal inference

When can causal effects be read off regression models in single time-point settings?

LINEAR MODEL with an interaction between A and W :

$$E(Y | A = a, W = w) = \beta_0 + \beta_1 a + \beta_2 w + \beta_3 aw$$

$$E(Y | A = 1, W = w) - E(Y | A = 0, W = w) = \beta_1 + \beta_3 w$$

$$E[E(Y | A = 1, W) - E(Y | A = 0, W)] = \beta_1 + \beta_3 E(W)$$

Regression coefficients generally cannot be interpreted as average treatment effects.

Failure of naive approach to causal inference

When can causal effects be read off regression models in single time-point settings?

GENERALIZED LINEAR MODEL (e.g., logistic model):


$$E(Y | A = a, W = w) = \text{expit}(\beta_0 + \beta_1 a + \beta_2 w)$$

$$E(Y | A = 1, W = w) - E(Y | A = 0, W = w) = \text{expit}(\beta_0 + \beta_1 + \beta_2 w) - \text{expit}(\beta_0 + \beta_2 w)$$

$$E[E(Y | A = 1, W) - E(Y | A = 0, W)] = E[\text{expit}(\beta_0 + \beta_1 + \beta_2 W) - \text{expit}(\beta_0 + \beta_2 W)]$$

Regression coefficients generally cannot be interpreted as average treatment effects.

Failure of naive approach to causal inference



$$\begin{aligned} Y \mid A_1 = a_1, L_1 = \ell_1, A_0 = a_0 &\sim \text{Normal}(1 + a_1 + 2\ell_1, 1) \\ A_1 \mid L_1 = \ell_1, A_0 = a_0 &\sim \text{Bernoulli}(\text{expit}(-1 + \ell_1)) \\ L_1 \mid A_0 = a_0 &\sim \text{Normal}(1 + a_0, 1) \\ A_0 &\sim \text{Bernoulli}(0.5) \end{aligned}$$

Failure of naive approach to causal inference

$$\begin{aligned}Y \mid A_1 = a_1, L_1 = \ell_1, A_0 = a_0 &\sim \text{Normal}(1 + a_1 + 2\ell_1, 1) \\A_1 \mid L_1 = \ell_1, A_0 = a_0 &\sim \text{Bernoulli}(\text{expit}(-1 + \ell_1)) \\L_1 \mid A_0 = a_0 &\sim \text{Normal}(1 + a_0, 1) \\A_0 &\sim \text{Bernoulli}(0.5)\end{aligned}$$

Using the G-computation formula, we can compute mean counterfactual outcomes corresponding to different treatment profiles:

$$\begin{aligned}E[Y(1, 1)] &= E[E[Y \mid A_1 = 1, L_1, A_0 = 1] \mid A_0 = 1] = E[2 + 2L_1 \mid A_0 = 1] = 6 \\E[Y(1, 0)] &= E[E[Y \mid A_1 = 0, L_1, A_0 = 1] \mid A_0 = 1] = E[1 + 2L_1 \mid A_0 = 1] = 5 \\E[Y(0, 1)] &= E[E[Y \mid A_1 = 1, L_1, A_0 = 0] \mid A_0 = 0] = E[2 + 2L_1 \mid A_0 = 0] = 4 \\E[Y(0, 0)] &= E[E[Y \mid A_1 = 0, L_1, A_0 = 0] \mid A_0 = 0] = E[1 + 2L_1 \mid A_0 = 0] = 3\end{aligned}$$

Failure of naive approach to causal inference

Can causal effects be read off the regression of Y on (A_1, L_1, A_0) ?

Effect of differing A_1 values but same A_0 value:

$$E[Y(1, 1) - Y(1, 0)] = 1 = E[Y | A_1 = 1, L_1, A_0 = 1] - E[Y | A_1 = 0, L_1, A_0 = 1]$$

$$E[Y(0, 1) - Y(0, 0)] = 1 = E[Y | A_1 = 1, L_1, A_0 = 0] - E[Y | A_1 = 0, L_1, A_0 = 0]$$

Effect of differing A_0 values but same A_1 value:

$$E[Y(1, 1) - Y(0, 1)] = 2 \neq 0 = E[Y | A_1 = 1, L_1, A_0 = 1] - E[Y | A_1 = 1, L_1, A_0 = 0]$$

$$E[Y(1, 0) - Y(0, 0)] = 2 \neq 0 = E[Y | A_1 = 0, L_1, A_0 = 1] - E[Y | A_1 = 0, L_1, A_0 = 0]$$


Effect of differing A_1 and A_0 values:

$$E[Y(1, 1) - Y(0, 0)] = 3 \neq 1 = E[Y | A_1 = 1, L_1, A_0 = 1] - E[Y | A_1 = 0, L_1, A_0 = 0]$$

$$E[Y(1, 0) - Y(0, 1)] = 1 \neq -1 = E[Y | A_1 = 0, L_1, A_0 = 1] - E[Y | A_1 = 1, L_1, A_0 = 0]$$

Failure of naive approach to causal inference

Why does this happen?




The regression of Y on (A_0, L_1, A_1) fixes L_1 : as such, the causal path between A_0 and Y is blocked. The observed (lack of) association between A_0 and Y is thus not causal.

This is a case of **treatment-confounder feedback**, whose presence often invalidates naive approaches to causal inference (see Chapter 20 of Hernán & Robins, 2018).

Causal methods are even more critical in the context of time-varying interventions!


Failure of naive approach to causal inference

Is treatment-confounder feedback present in these examples?



Failure of naive approach to causal inference

Is treatment-confounder feedback present in these examples?



Simple estimation procedures

The IPTW identification formula (slide 13) suggests a simple estimation strategy.

First, we may construct estimates $g_{0n}, g_{1n}, \dots, g_{Kn}$ of propensity scores g_0, g_1, \dots, g_K using our favorite regression estimator for binary outcomes:

- $g_0(\ell_0)$: mean of A_0 given $L_0 = \ell_0$;
- $g_1(\bar{\ell}_1)$: mean of A_1 given $\bar{L}_1 = \bar{\ell}_1$ and $A_0 = 1$;
- ...
- $g_K(\bar{\ell}_K)$: mean of A_K given $\bar{L}_K = \bar{\ell}_K$ and $\bar{A}_{K-1} = 1_K$.

Then, we can compute the corresponding **IPTW estimator** of $E[Y(1, 1, \dots, 1)]$ as

$$\psi_{n, IPTW} := \frac{1}{n} \sum_{i=1}^n \left\{ \frac{A_{0i} A_{1i} \dots A_{Ki}}{g_{0n}(L_{0i}) g_{1n}(\bar{L}_{1i}) \dots g_{Kn}(\bar{L}_{Ki})} \right\} Y_i .$$

Simple estimation procedures

The G-computation formula suggests another natural estimation strategy.

We first illustrate in the three time-point setting, in which case the estimand is simply

$$E [E [E (Y | A_1 = 1, L_1, A_0 = 1, L_0) | A_0 = 1, L_0]] .$$

- 1 Get estimate $\bar{Q}_{2,n}$ of the regression function

$$\bar{Q}_2(\ell_1, \ell_0) := E (Y | A_1 = 1, L_1 = \ell_1, A_0 = 1, L_0 = \ell_0) ,$$

taking Y as outcome and (A_1, L_1, A_0, L_0) as covariate vector.

- 2 Get estimate F_{1n} of the distribution F_1 of L_1 given $A_0 = 1$ and L_0 , taking L_1 as outcome and (A_0, L_0) as covariate vector, and average \bar{Q}_{2n} over F_{1n} to get

$$\bar{Q}_{1n}(\ell_0) := E_{F_{1n}} [\bar{Q}_{2n}(L_1, \ell_0) | A_0 = 1, L_0 = \ell_0] = \int \bar{Q}_{2n}(\ell_1, \ell_0) dF_{1n}(\ell_1 | \ell_0) .$$

- 3 Average \bar{Q}_{1n} over the empirical distribution F_{0n} of L_0 to get

$$\psi_{G,n} := \bar{Q}_{0n} := E_{F_{0n}} [\bar{Q}_{1n}(L_0)] = \frac{1}{n} \sum_{i=1}^n \bar{Q}_{1n}(L_{0i}) .$$

Simple estimation procedures

In general, STEP 2 must be performed $K - 1$ times, each time taking the output from the previous step and averaging it out relative to an estimated conditional distribution.

This can be challenging since it requires the user to:

- estimate an entire conditional distribution;
- average out an estimated regression function with respect to this estimate.

Instead, a simpler **sequential regression** approach can be used (Bang & Robins, 2005).

To obtain \bar{Q}_{1n} , we could run *our favorite regression method* on all observations with $A_0 = 1$ using $\bar{Q}_{2n}(L_{1i}, L_{0i})$ as outcome and L_{0i} as covariate.

Simple estimation procedures

Target estimand: $E[E[E(Y | A_1 = 1, L_1, A_0 = 1, L_0) | A_0 = 1, L_0]]$

Y	A_1	L_1	A_0	L_0
0	1	0.7	1	2.1
1	1	-0.2	0	-1.6
1	1	2.0	1	0.3
0	0	6.9	1	1.4
1	0	3.1	0	0.9
0	1	-5.2	1	-3.1
1	0	5.2	1	2.5
0	1	-1.1	1	-0.1

Simple estimation procedures

Target estimand: $E[E[E(Y | A_1 = 1, L_1, A_0 = 1, L_0) | A_0 = 1, L_0]]$

Y	A ₁	L ₁	A ₀	L ₀
0	1	0.7	1	2.1
1	1	-0.2	0	-1.6
1	1	2.0	1	0.3
0	0	6.9	1	1.4
1	0	3.1	0	0.9
0	1	-5.2	1	-3.1
1	0	5.2	1	2.5
0	1	-1.1	1	-0.1

STEP 1: Regress Y on L₁ and L₀ among those with A₁ = A₀ = 1. → \bar{Q}_{2n}

Simple estimation procedures

Target estimand: $E[E[E(Y | A_1 = 1, L_1, A_0 = 1, L_0) | A_0 = 1, L_0]]$

Y	A ₁	L ₁	A ₀	L ₀	$\bar{Q}_{2n}(L_1, L_0)$
0	1	0.7	1	2.1	0.34
1	1	-0.2	0	-1.6	0.19
1	1	2.0	1	0.3	0.26
0	0	6.9	1	1.4	0.31
1	0	3.1	0	0.9	0.29
0	1	-5.2	1	-3.1	0.15
1	0	5.2	1	2.5	0.36
0	1	-1.1	1	-0.1	0.25

STEP 1: Regress Y on L₁ and L₀ among those with A₁ = A₀ = 1. → \bar{Q}_{2n}

Compute $\bar{Q}_{2n}(L_1, L_0)$ for every patient.

Simple estimation procedures

Target estimand: $E[E[E(Y | A_1 = 1, L_1, A_0 = 1, L_0) | A_0 = 1, L_0]]$

Y	A ₁	L ₁	A ₀	L ₀	$\bar{Q}_{2n}(L_1, L_0)$
0	1	0.7	1	2.1	0.34
1	1	-0.2	0	-1.6	0.19
1	1	2.0	1	0.3	0.26
0	0	6.9	1	1.4	0.31
1	0	3.1	0	0.9	0.29
0	1	-5.2	1	-3.1	0.15
1	0	5.2	1	2.5	0.36
0	1	-1.1	1	-0.1	0.25

STEP 2: Regress $\bar{Q}_{2n}(L_1, L_0)$ on L_0 among those with $A_0 = 1$. $\rightarrow \bar{Q}_{1n}$

Simple estimation procedures

Target estimand: $E[E[E(Y | A_1 = 1, L_1, A_0 = 1, L_0) | A_0 = 1, L_0]]$

Y	A ₁	L ₁	A ₀	L ₀	$\bar{Q}_{2n}(L_1, L_0)$	$\bar{Q}_{1n}(L_0)$
0	1	0.7	1	2.1	0.34	0.40
1	1	-0.2	0	-1.6	0.19	0.22
1	1	2.0	1	0.3	0.26	0.20
0	0	6.9	1	1.4	0.31	0.36
1	0	3.1	0	0.9	0.29	0.21
0	1	-5.2	1	-3.1	0.15	0.34
1	0	5.2	1	2.5	0.36	0.20
0	1	-1.1	1	-0.1	0.25	0.29

STEP 2: Regress $\bar{Q}_{2n}(L_1, L_0)$ on L_0 among those with $A_0 = 1$. $\rightarrow \bar{Q}_{1n}$
Compute $\bar{Q}_{1n}(L_0)$ for every patient.

Simple estimation procedures

Target estimand: $E[E[E(Y | A_1 = 1, L_1, A_0 = 1, L_0) | A_0 = 1, L_0]]$

Y	A ₁	L ₁	A ₀	L ₀	$\bar{Q}_{2n}(L_1, L_0)$	$\bar{Q}_{1n}(L_0)$
0	1	0.7	1	2.1	0.34	0.40
1	1	-0.2	0	-1.6	0.19	0.22
1	1	2.0	1	0.3	0.26	0.20
0	0	6.9	1	1.4	0.31	0.36
1	0	3.1	0	0.9	0.29	0.21
0	1	-5.2	1	-3.1	0.15	0.34
1	0	5.2	1	2.5	0.36	0.20
0	1	-1.1	1	-0.1	0.25	0.29

STEP 3: Average out values of $\bar{Q}_{1n}(L_0)$ over all patients to get $\psi_{n,GCOMP} := \bar{Q}_{0n}$.

Improved estimation procedures

Much like in the single time-point case, a hybrid between the G-computation and IPTW estimators can be constructed, and enjoys improved properties.

The **augmented IPTW (AIPTW) estimator** is given by

$$\psi_{n,AIPTW} := \bar{Q}_{0n} + \frac{1}{n} \sum_{i=1}^n \sum_{j=0}^K \left\{ \frac{A_{0i} A_{1i} \dots A_{ji}}{\bar{g}_{jn}(\bar{L}_{ji})} \right\} \left[\bar{Q}_{(j+1)n}(\bar{L}_{(j+1)i}) - \bar{Q}_{jn}(\bar{L}_{ji}) \right],$$

where \bar{Q}_{0n} is simply the G-computation estimator (see slide 23) (Robins ref).

Since it **builds upon estimates of all outcome regressions and propensity scores**, the construction of this estimator requires more effort than for estimators seen so far.

However, this estimator **enjoys double-robustness**, and can be used to construct **valid confidence intervals**, even when flexible learning strategies (e.g., Super Learner) are used to estimate the outcome regressions and propensity scores.

Improved estimation procedures

What does **double robustness** refer to in the context of multi time-point interventions?

The estimator $\psi_{n,AIPW}$ built upon estimators

$$\bar{Q}_n := (\bar{Q}_{1n}, \bar{Q}_{2n}, \dots, \bar{Q}_{(K+1)n}) \text{ and } g_n := (g_{0n}, g_{1n}, \dots, g_{Kn})$$

is doubly-robust, in the sense that it is consistent (i.e., hits the target) **provided either \bar{Q}_n hits the target \bar{Q}_0 or g_n hits the target g_0 .**

Scenario 1: \bar{Q}_n is consistent

	time 0	time 1	time 2	time 3	...	end
\bar{Q}	—	✓	✓	✓	...	✓
g						

Improved estimation procedures

What does **double robustness** refer to in the context of multi time-point interventions?

The estimator $\psi_{n,AIPTW}$ built upon estimators

$$\bar{Q}_n := (\bar{Q}_{1n}, \bar{Q}_{2n}, \dots, \bar{Q}_{(K+1)n}) \text{ and } g_n := (g_{0n}, g_{1n}, \dots, g_{Kn})$$

is doubly-robust, in the sense that it is consistent (i.e., hits the target) **provided either \bar{Q}_n hits the target \bar{Q}_0 or g_n hits the target g_0 .**

Scenario 2: g_n is consistent

	time 0	time 1	time 2	time 3	...	end
\bar{Q}					...	
g	✓	✓	✓	✓	...	✓

Improved estimation procedures

What does **double robustness** refer to in the context of multi time-point interventions?

The estimator $\psi_{n,AIPTW}$ built upon estimators

$$\bar{Q}_n := (\bar{Q}_{1n}, \bar{Q}_{2n}, \dots, \bar{Q}_{(K+1)n}) \text{ and } g_n := (g_{0n}, g_{1n}, \dots, g_{Kn})$$

is doubly-robust, in the sense that it is consistent (i.e., hits the target) **provided either \bar{Q}_n hits the target \bar{Q}_0 or g_n hits the target g_0 .**

Scenario 3: g_n is consistent at early times, \bar{Q}_n is consistent at later times

	time 0	time 1	time 2	time 3	...	end
\bar{Q}			✓	✓	...	✓
g	✓	✓				

Improved estimation procedures

The targeted maximum likelihood estimation (TMLE) framework provides a recipe for constructing a G-computation estimator that is also doubly-robust.

(Bang & Robins 2005; van der Laan & Gruber, 2012)

This estimator is constructed like the 'sequential regression' form of the G-computation estimator, but includes a refinement step after each \bar{Q} estimator is obtained.

Using a given estimator g_n of g_0 , the algorithm proceeds as follows:

- (1) build estimate $\bar{Q}_{(K+1)n}$ of \bar{Q}_{K+1} using your favorite regression tool;
- (2) refine $\bar{Q}_{(K+1)n}$, say to $\bar{Q}_{(K+1)n}^*$, to make it a targeted estimate; (★)
- (3) build estimate \bar{Q}_{Kn} of \bar{Q}_K by regressing $\bar{Q}_{(K+1)n}^*(\bar{L}_K)$ using your favorite regression tool;
- (4) refine \bar{Q}_{Kn} , say to \bar{Q}_{Kn}^* , to make it a targeted estimate; (★)
- ...
- (2K+1) build estimate \bar{Q}_{1n} of \bar{Q}_1 by regressing \bar{Q}_{2n}^* using your favorite regression tool;
- (2K+2) refine \bar{Q}_{1n} , say to \bar{Q}_{1n}^* , to make it a targeted estimate; (★)
- (2K+3) take final estimate to be the G-computation estimator $\psi_{n, TMLE} := \bar{Q}_{0n}^* := \frac{1}{n} \sum_{i=1}^n \bar{Q}_{1n}^*(L_{0i})$.

What do these so-called **targeting steps** consist of?

STEP (2): Get slope estimate β_n from logistic regression with outcome Y , single covariate

$$Z := (A_0 A_1 \dots A_K) / \bar{g}_{Kn}(\bar{L}_K)$$

and offset term $\text{logit } \bar{Q}_{(K+1)n}(\bar{L}_K)$ using only data with $A_0 = A_1 = \dots = A_K = 1$.

Set $\bar{Q}_{(K+1)n}^* := \text{expit}[\text{logit } \bar{Q}_{(K+1)n} + \beta_n / \bar{g}_{Kn}]$.

STEP (4): Get slope estimate β_n from logistic regression with outcome $\bar{Q}_{(K+1)n}^*$, single covariate

$$Z := (A_0 A_1 \dots A_{K-1}) / \bar{g}_{(K-1)n}(\bar{L}_{K-1})$$

and offset term $\text{logit } \bar{Q}_{Kn}(\bar{L}_{K-1})$ using only data with $A_0 = A_1 = \dots = A_{K-1} = 1$.

Set $\bar{Q}_{Kn}^* := \text{expit}[\text{logit } \bar{Q}_{Kn} + \beta_n / \bar{g}_{(K-1)n}]$.

STEP (2K+2): Get slope estimate β_n from logistic regression with outcome \bar{Q}_{2n}^* , single covariate

$$Z := A_0 / \bar{g}_{0n}(\bar{L}_0)$$

and offset term $\text{logit } \bar{Q}_{1n}(\bar{L}_0)$ using only data with $A_0 = 1$. Set $\bar{Q}_{1n}^* := \text{expit}[\text{logit } \bar{Q}_{1n} + \beta_n / \bar{g}_{0n}]$.

Improved estimation procedures

Properties of estimation procedures outlined

	difficulty	$\bar{Q} + \bar{g}$		$\bar{Q} + \bar{g}$		$\bar{Q} + \bar{g}$	
		target	ci	target	ci	target	ci
IPTW	+			✓		✓	
G-COMP	++	✓				✓	
AIPTW	+++	✓		✓		✓	✓
TMLE	++++	✓	✓	✓	✓	✓	✓

$\bar{Q} + \bar{g}$: outcome regressions estimated well but not propensity scores

$\bar{Q} + \bar{g}$: propensity scores estimated well but not outcome regressions

$\bar{Q} + \bar{g}$: outcome regressions and propensity scores estimated well

target : does the estimator hit the right target?

ci : is valid inference possible and readily available, even when flexible learning strategies (such as Super Learner) are used?

There is substantial work underway to produce novel estimators with even better properties. All of these innovations are based on the idea of TMLE.

■ **Enhanced robustness**

- Consistent estimation is possible under a wider range of scenarios than depicted on slides 27, 28 and 29.
- For this, more complicated procedures are needed and are being developed. (Luedtke et al. 2017; Rotnitzky et al., 2017)

■ **Robust inference**

- Typical double-robustness only refers to consistency.
- However, constructing doubly-robust CI and tests is a much more important task. It is also very difficult in multi time-point settings. (Benkeser et al., 2017)

■ **Targeted estimation of propensity scores**

- Particularly when there are many potential confounders, good performance may be difficult to achieve in smaller samples using the methods described so far.
- Collaborative TMLE allows a smarter, data-driven selection of propensity score estimators to improve performance in such cases. (van der Laan & Gruber, 2010)

Devenir Après Interruption de la FIV (DAIFI) study

(Chapter 25 of van der Laan & Rose, 2011)

- > **Motivating question:** How successful is IVF therapy in France?
- > **Study sample:** All women who received a first IVF cycle at two French IVF units between 1998 and 2002 and were under 42 years of age at initiation.
- > **Intervention considered:** Four successive IVF cycles.
- > **Outcome:** Successful delivery arising from IVF.
- > **Observational challenges:**
 - Some couples abandon mid-course without a successful delivery.
 - Common factors likely influence discontinuation and overall chance of success.
- > **Observed probability of success:**

1st cycle	2nd cycle	3rd cycle	4th cycle
22%	32%	35%	37%

Devenir Après Interruption de la FIV (DAIFI) study

(Chapter 25 of van der Laan & Rose, 2011)

> **Data structure:**

L_0 = information recorded at baseline (first cycle)
(e.g., age, center, # of embryos transferred or frozen, success of first cycle)

A_0 = second cycle attempted

L_1 = success of first two cycles

A_1 = third cycle attempted

L_2 = success of first three cycles

A_2 = fourth cycle attempted

Y = success of first four cycles

> **Causal estimand of interest:** $E[Y(1, 1, 1)]$

> **Result of TMLE analysis:** 50.5% (95%CI: 48.0–53.0)


HVTN 505 study

(Chapter 11 of van der Laan & Rose, 2018)

- > **Motivating question:** Does novel HIV vaccine exhibit sieve effects?
- > **Study sample:** Phase II preventive HIV vaccine efficacy trial. 2,405 participants were randomized 1:1 to receive candidate vaccine or placebo.
- > **Interventions considered:** (Active vs control vaccine) + (no loss to follow-up).
- > **Outcome:** Infection with specific genotypes of HIV.
- > **Observational challenges:**
 - Participants may have unblinded using home HIV tests – higher dropout in control arm.
 - Risk behaviors may be informative of participant dropout and HIV infection risk.


HVTN 505 study

(Chapter 11 of van der Laan & Rose, 2018)



HVTN 505 study

(Chapter 11 of van der Laan & Rose, 2018)



Positivity violations are not uncommon, especially in multi time-point settings.
(Petersen, 2012)

- They may happen **in a study sample by chance** (i.e., practical violations).
 - Can use estimators that borrow more information from other patients.
- They may instead be a **fundamental feature** of the particular combination of population and treatment considered.
 - Can **change the reference population** to exclude subgroups for which the average treatment effect cannot be learned.
 - Can instead **focus on more realistic interventions** for which positivity holds, such as dynamic treatments reflecting clinical practice.

Dynamic treatment rules

Counterfactuals defined by fixed treatment profiles are often neither particularly clinically interesting nor supported by data.

Treatment decisions are usually dynamic and incorporate real-time patient information.

Example: mercaptopurine in IBD patients

- static intervention: 'always treat' versus 'never treat'
- if patient develops signs of liver damage, therapy is usually stopped
- liver function is a time-varying confounder between treatment status and survival
- if poor liver function is a contraindication for therapy, it may not be possible to observe treatment adherence among patients with recent liver failure
- static intervention is unrealistic and not identifiable
- dynamic intervention: 'treat while liver function permits it' versus 'never treat'

$$d(t) = \begin{cases} 1 & : \text{if recent liver function is adequate} \\ 0 & : \text{otherwise} \end{cases}$$

Dynamic treatment rules

Counterfactuals can be naturally defined in terms of **dynamic treatment rules encoding treatment decisions that possibly depend on current and past patient info.**

In the mercaptopurine example, we may want to learn about the average effect

$$ATE(d, d_0) := E[Y(d)] - E[Y(d_0)]$$

of rule d enforcing treatment whenever liver function permits it and rule d_0 enforcing no mercaptopurine use.

All methods discussed so far can be adapted for use with dynamic treatment rules. Wherever we imposed $\bar{A}_k = \bar{a}_k$, we now instead set \bar{A}_k to equal $d(\bar{L}_k)$ or $d_0(\bar{L}_k)$.

Dynamic treatment rules

A vast subfield of causal inference focuses on the quest for optimal rules.
(Chakraborty & Moodie, 2013)

Of all candidate treatment rules, which one results in the most beneficial average treatment effect?

This is referred to as an **optimal dynamic treatment regime**.

Finding the optimal rule and constructing confidence intervals for the average effect corresponding to this rule using the same data is challenging.

(see, e.g.: Laber et al., 2014; Luedkte & van der Laan, 2016)

Key points of Chapter 5

- Methods for time-varying interventions are extremely versatile, and can be used to tackle loss to follow-up and missing data.
- The G-computation and IPTW formulas can be extended to time-varying settings.
- Standard regression should never be used to study time-varying interventions.
- While model-based G-computation and IPTW estimators are still available, matching is no longer an option at all.
- Doubly-robust estimators should be preferred as they confer efficiency, additional robustness and the ability to use flexible nuisance estimators.
- Much work is currently being done to further improve doubly-robust estimators.
- Dynamic treatment rules may better reflect realistic interventions and prevent positivity violations.
- Methods above can be used to estimate the average effect of dynamic rules.
- Identifying and making inference about optimal dynamic rules is more difficult.

References and additional reading

References:

Bang H, Robins JM (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics*; 61(4)962-973. doi: [10.1111/j.1541-0420.2005.00377.x/](https://doi.org/10.1111/j.1541-0420.2005.00377.x/)

Benkeser, D, Carone, M, van der Laan, MJ & Gilbert, P (2017). Doubly-robust nonparametric inference on the average treatment effect. *Biometrika*; 104(4)863-880. doi: [10.1093/biomet/asx053](https://doi.org/10.1093/biomet/asx053).

Chakraborty B, Moodie EE (2013). *Statistical Methods for Dynamic Treatment Regimes*. Springer New York.

Hernán MA, Robins JM (2018). *Causal Inference*. Chapman & Hall/CRC. *Forthcoming – draft available online*.

Laber EB, Lizotte DJ, Qian M, Pelham WE, Murphy SA (2014). Dynamic treatment regimes: technical challenges and applications. *Electronic Journal of Statistics*; 8(1)1225-1272. doi: [10.1214/14-EJS920](https://doi.org/10.1214/14-EJS920)

Luedtke AR, van der Laan MJ (2016). Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. *Annals of Statistics*; 44(2)713-742. doi: [10.1214/15-AOS1384](https://doi.org/10.1214/15-AOS1384)

Luedtke AR, Sofrygin O, van der Laan MJ, Carone M (2017). Sequential double-robustness in right-censored longitudinal models. *ArXiv technical report*. ArXiv: [1705.02459](https://arxiv.org/abs/1705.02459)

Petersen ML, Porter KE, Gruber S, Wang Y, van der Laan MJ (2012). Diagnosis and responding to violations in the positivity assumption. *Statistical Methods in Medical Research*; 21(1)31-54. doi: [10.1177/0962280210386207](https://doi.org/10.1177/0962280210386207)

References and additional reading

References:

Robins JM (1986). A new approach to causal inference in mortality studies with a sustained exposure period – application to control of the healthy worker survivor effect. *Mathematical Modelling*, 9(7)1393-1512. doi: [10.1016/0270-0255\(86\)90088-6](https://doi.org/10.1016/0270-0255(86)90088-6).

Rotnitzky A, Robins JM, Babino L (2017). On the multiply robust estimation of the mean of the g-functional. *ArXiv technical report*. ArXiv: [1705.08582](https://arxiv.org/abs/1705.08582)

van der Laan MJ, Gruber S (2010). Collaborative double robust targeted maximum likelihood estimation. *International Journal of Biostatistics*; 6(1). doi: [10.2202/1557-4679.1182](https://doi.org/10.2202/1557-4679.1182)

van der Laan MJ, Rose S (2011). *Targeted Learning: Causal Inference for Observation and Experimental Data*. Springer New York. [10.1007/978-1-4419-9782-1](https://doi.org/10.1007/978-1-4419-9782-1)

van der Laan MJ, Gruber S (2012). Targeted minimum loss based estimation of causal effects of multiple time point interventions. *International Journal of Biostatistics*; 8(1). doi: [10.1515/1557-4679.1370](https://doi.org/10.1515/1557-4679.1370)

van der Laan MJ, Rose S (2018). *Targeted Learning in Data Science: Causal Inference for Complex Longitudinal Studies*. Springer New York. [10.1007/978-3-319-65304-4](https://doi.org/10.1007/978-3-319-65304-4)