

Chapter 19: Time-Varying Treatments

Contents

1	19.1 The Causal Effect of Time-Varying Treatments (p. 255)	1
1.1	The Causal Estimand	2
2	19.2 Treatment Strategies (p. 256)	2
3	19.3 Sequentially Randomized Experiments (p. 257)	3
3.1	Connection to the G-Formula	4
4	19.4 Sequential Exchangeability (p. 259)	4
4.1	Comparison with Time-Fixed Exchangeability	4
4.2	Sequential Positivity	5
5	19.5 Identifiability Under Some but Not All Treatment Strategies (p. 261)	5
5.1	Implications for Study Design	5
6	19.6 Time-Varying Confounding and Time-Varying Confounders (p. 265)	6
6.1	The Feedback Loop	6
7	Summary	7
8	References	7

In the earlier chapters of this book we examined methods for estimating the causal effect of a **time-fixed** treatment — one that is assigned once at the start of follow-up and then held constant throughout the study period. In practice, many interventions and exposures unfold over time: a patient may initiate, discontinue, and restart a medication; a worker may change jobs or occupational exposures from year to year; a diet may evolve across decades of life. Part III of the textbook extends the causal inference framework to these **time-varying treatments**.

This chapter is based on Hernán and Robins (2020, chap. 19, pp. 255–265).

Central challenge: When treatments vary over time, the causal question becomes richer — we ask not just “what if treated vs. untreated?” but “what if treated according to *this* sequence of decisions?” Sequential randomization is the conceptual gold standard, and sequential exchangeability is the identifying assumption that allows observational data to stand in for it.

1 19.1 The Causal Effect of Time-Varying Treatments (p. 255)

Let time be indexed by $k = 0, 1, \dots, K$, where $k = 0$ is the baseline and K is the end of the study period. The **time-varying treatment** at time k is denoted A_k , which may be binary (e.g., $A_k = 1$ if the individual is on treatment at time k , 0 otherwise) or take more general values.

Definition 1.1 (Treatment History). The observed treatment **history** up to time k is the vector

$$\bar{A}_k = (A_0, A_1, \dots, A_k).$$

Definition 1.2 (Covariate History). The **covariate history** up to time k is

$$\bar{L}_k = (L_0, L_1, \dots, L_k),$$

where L_k represents all measured covariates recorded at time k (e.g., clinical measurements, lab values, questionnaire responses).

The outcome Y is measured at the end of follow-up, after A_K . A common data structure in person-time format is therefore:

$$L_0 \rightarrow A_0 \rightarrow L_1 \rightarrow A_1 \rightarrow \dots \rightarrow L_K \rightarrow A_K \rightarrow Y.$$

Definition 1.3 (Potential Outcome Under a Treatment History). For a fully specified treatment history $\bar{a}_K = (a_0, a_1, \dots, a_K)$, the **potential outcome** $Y^{\bar{a}_K}$ is the outcome that would have been observed had the individual's treatment been set to $\bar{a}_K$ throughout the follow-up, regardless of what treatment they actually received.

When $K = 0$ this reduces to the familiar time-fixed potential outcome Y^a . With K time points, the number of distinct potential outcomes grows exponentially with K , which motivates the use of treatment **strategies** rather than enumerating all histories.

Notation note: Some texts use $Y^{\bar{a}}$ or $Y(\bar{a})$ for the same quantity. We follow Hernán and Robins and write $Y^{\bar{a}_K}$ when the full treatment history up to K matters, and Y^g when the treatment is assigned according to a strategy g (defined in Section 19.2).

The counterfactual world framework still applies: for each individual there exists a potential outcome $Y^{\bar{a}_K}$ for every conceivable treatment history, but we observe only the one that corresponds to the treatment the individual actually received.

1.1 The Causal Estimand

The **average causal effect** of treatment history $\bar{a}_K$ versus $\bar{a}_K^*$ is

$$E[Y^{\bar{a}_K}] - E[Y^{\bar{a}_K^*}].$$

When the treatment is binary at each time point, there are 2^{K+1} possible histories, so comparing all pairs quickly becomes unwieldy. In most applications we focus on a small number of pre-specified **treatment strategies**, described in the next section.

2 19.2 Treatment Strategies (p. 256)

A **treatment strategy** (also called a **treatment regime** or **dynamic treatment regime**) is a rule g that specifies, at each time k , what treatment value a_k should be assigned, possibly as a function of the individual's past treatment and covariate history.

Definition 2.1 (Treatment Strategy). A treatment strategy $g = (g_0, g_1, \dots, g_K)$ is a collection of decision functions where each g_k maps the observed history $(\bar{a}_{k-1}, \bar{l}_k)$ to a treatment value:

$$g_k : (\bar{a}_{k-1}, \bar{l}_k) \mapsto a_k.$$

The counterfactual outcome under strategy g is denoted Y^g .

Definition 2.2 (Static Strategies). A **static strategy** assigns the same treatment value at every time point regardless of the evolving covariate history.

Examples:

- “Always treat”: $g_k(\bar{a}_{k-1}, \bar{l}_k) = 1$ for all k .
- “Never treat”: $g_k(\bar{a}_{k-1}, \bar{l}_k) = 0$ for all k .
- “Treat at time 0 only”: $g_0 = 1, g_k = 0$ for $k \geq 1$.

The potential outcome under “always treat” is $Y^{\bar{a}_K=(1,1,\dots,1)}$, often written $Y^{\bar{1}}$, and under “never treat” it is $Y^{\bar{0}}$.

Definition 2.3 (Dynamic Strategies). A **dynamic strategy** adapts treatment decisions based on the individual’s evolving history. The counterfactual Y^g under a dynamic strategy g is defined analogously to the static case — it is the outcome that would have been observed if, at every time k , treatment had been set according to g_k applied to the individual’s *counterfactual* covariate history under g .

For example:

- “Treat whenever the CD4 count falls below 200 cells/mm³.”
- “Initiate therapy if blood pressure exceeds 140 mmHg, then continue indefinitely.”

Dynamic strategies are clinically natural: many real-world treatment guidelines are themselves dynamic.

Threshold strategies: A common subclass of dynamic strategies are “threshold” or “if–then” rules, in which treatment is initiated or maintained once a covariate crosses a threshold. These arise naturally in comparative effectiveness research (e.g., “start antiretrovirals when CD4 < 350 vs. CD4 < 500”).

Adaptive strategies: Some dynamic strategies use the treatment history itself. For example, “if you were on treatment last period and your disease is controlled, continue; otherwise switch to an alternative regimen.” These are called **adaptive treatment strategies** or **multi-stage treatment strategies** and are the focus of SMART (sequential multiple assignment randomized trial) designs.

3 19.3 Sequentially Randomized Experiments (p. 257)

The conceptual benchmark for estimating the effects of time-varying treatments is the **sequentially randomized experiment**.

Definition 3.1 (Sequentially Randomized Experiment). A sequentially randomized experiment is a longitudinal study in which, at each time k , the treatment A_k is randomly assigned with probabilities that may depend on the observed past $(\bar{A}_{k-1}, \bar{L}_k)$, but not on any unmeasured variables.

In a sequentially randomized experiment, at each time k the treatment probability satisfies

$$\Pr[A_k = a_k \mid \bar{A}_{k-1} = \bar{a}_{k-1}, \bar{L}_k = \bar{l}_k, U] = \Pr[A_k = a_k \mid \bar{A}_{k-1} = \bar{a}_{k-1}, \bar{L}_k = \bar{l}_k]$$

for all values of unmeasured variables U . This is the time-varying analogue of single-time-point randomization: within each “stratum” defined by $(\bar{A}_{k-1}, \bar{L}_k)$, treatment at time k is effectively randomized.

3.1 Connection to the G-Formula

In a sequentially randomized experiment, the counterfactual mean under strategy g equals

$$E[Y^g] = \sum_{\bar{l}_K} E[Y \mid \bar{A}_K = g(\bar{l}_K), \bar{L}_K = \bar{l}_K] \prod_{k=0}^K f(l_k \mid \bar{A}_{k-1} = g_{k-1}(\bar{l}_{k-1}), \bar{L}_{k-1} = \bar{l}_{k-1}),$$

an expression that we will expand in Chapter 21. For now, the key point is that **sequential randomization ensures identification**: the counterfactual mean can be expressed as a function of observable data.

Why sequential randomization differs from a single randomization:

In a standard (single-time-point) randomized trial, randomization happens once at baseline. This guarantees marginal exchangeability: $Y^a \perp\!\!\!\perp A$. But if new covariates L_1 emerge after baseline and influence both A_1 and Y , we need randomization at time 1 as well. In a sequentially randomized experiment, this is guaranteed at every time point, conditionally on the covariate history.

Practical example: SMART designs in clinical trials operationalize sequential randomization. Participants are randomized to an initial treatment, then re-randomized at a pre-specified decision point (often based on response status). This design yields valid inference for dynamic strategies under clearly stated assumptions.

4 19.4 Sequential Exchangeability (p. 259)

In observational studies we cannot guarantee randomization at any time point. The identifying assumption that replaces sequential randomization is **sequential (or time-varying) exchangeability**.

Definition 4.1 (Sequential Exchangeability). Sequential exchangeability holds if, for every time k and every strategy g ,

$$Y^g \perp\!\!\!\perp A_k \mid \bar{A}_{k-1} = \bar{g}_{k-1}(\bar{L}_{k-1}), \bar{L}_k,$$

where $\bar{g}_{k-1}$ denotes the treatment history up to $k-1$ as dictated by g .

Intuitively, this condition says: *among individuals who followed strategy g up to time $k-1$ and share the same covariate history $\bar{L}_k$, the treatment actually received at time k is independent of all potential outcomes*. In other words, all variables that jointly predict A_k and are prognostic for Y^g have been measured in $\bar{L}_k$.

4.1 Comparison with Time-Fixed Exchangeability

In the time-fixed setting, exchangeability reads $Y^a \perp\!\!\!\perp A \mid L$: potential outcomes are independent of treatment, given measured covariates. Sequential exchangeability extends this to every time point in sequence.

A sufficient (but not necessary) condition for sequential exchangeability is that $\bar{L}_k$ contains all **common causes** of A_k and Y^g that are not caused by past treatment. This is often formalized using a DAG: sequential exchangeability holds if all backdoor paths from A_k to Y that are not mediated by past treatment are blocked by $\bar{L}_k$.

Why “sequential”? The adjective emphasizes that the condition must hold at *each* time k in the sequence, not just overall. A single global conditional independence statement is insufficient because new confounders may enter after baseline, and past treatment may affect future confounders (see Chapter 20).

Unmeasured confounders: As in the time-fixed case, sequential exchangeability fails when there are unmeasured common causes of A_k and Y that are not blocked by $\bar{L}_k$. Part III methods (g-formula, IP weighting, g-estimation) all require this assumption.

Sensitivity analysis: Because sequential exchangeability is untestable from the data alone, sensitivity analysis—asking how strong unmeasured confounding would need to be to explain away the result—is even more important in longitudinal settings.

4.2 Sequential Positivity

In addition to sequential exchangeability, identification requires **sequential positivity**: for each time k and every combination of $(\bar{a}_{k-1}, \bar{l}_k)$ that can occur under strategy g ,

$$\Pr[A_k = g_k(\bar{a}_{k-1}, \bar{l}_k) \mid \bar{A}_{k-1} = \bar{a}_{k-1}, \bar{L}_k = \bar{l}_k] > 0.$$

This says that the treatment mandated by strategy g at time k must be possible for every subgroup defined by the observed history. Violations of positivity are common in longitudinal data when, for example, clinical guidelines forbid certain treatment sequences, or when patients with particular covariate histories are never observed to receive a given treatment.

5 19.5 Identifiability Under Some but Not All Treatment Strategies (p. 261)

An important subtlety in the time-varying setting is that sequential exchangeability and positivity may hold for some strategies but not for others.

Consider a study in which patients are treated based on their measured covariate history. The strategy “treat if $L_k > c$ ” respects the observed treatment assignment mechanism, so positivity holds by construction. By contrast, the strategy “treat if $L_k > c$ *and* the patient’s unmeasured frailty is low” cannot be evaluated because the unmeasured component violates exchangeability.

Definition 5.1 (Identifiable Strategy). A strategy g is **identifiable** from the observed data if:

1. Sequential exchangeability holds for g at every time k .
2. Sequential positivity holds for g at every time k .
3. The consistency assumption holds: $Y = Y^g$ whenever $\bar{A}_K = g(\bar{L}_K)$.

When only a subset of strategies is identifiable, we restrict inference to that subset and explicitly acknowledge which strategies cannot be compared from the available data.

5.1 Implications for Study Design

This identifiability result has direct implications for how studies should be designed and how results should be interpreted:

- Define the target strategy *before* data collection. If the strategy requires data that were not collected (e.g., a biomarker not measured), it cannot be emulated.
- Report the range of strategies for which the analysis is valid. Extrapolating to “always treat” when no individuals were observed on that strategy throughout follow-up requires strong positivity assumptions.
- Consider **grace periods** in strategy definitions. Requiring exact adherence to “treat every month if $L_k > c$ ” may be so strict that few observed individuals satisfy it. Allowing a grace period of ± 1 month can improve positivity at the cost of defining a slightly different (but more realistic) strategy.

Structural vs. random violations of positivity:

A *structural* violation arises when it is literally impossible for an individual with covariate history $\bar{l}_k$ to receive treatment a_k under strategy g . For example, a patient who has died cannot receive further treatment. This is not a problem if the strategy is defined to handle such situations (e.g., “treat until death”).

A *random* violation arises when, due to chance or small sample size, no observed individuals in stratum $(\bar{a}_{k-1}, \bar{l}_k)$ received treatment a_k even though it was in principle possible. Random violations are less severe but still require careful handling, particularly when using parametric models for IP weighting.

6 19.6 Time-Varying Confounding and Time-Varying Confounders (p. 265)

A central complication in the analysis of time-varying treatments is the presence of **time-varying confounders** — covariates that change over time and that are **causally affected by past treatment** while simultaneously confounding future treatment assignments.

Definition 6.1 (Time-Varying Confounder). A variable L_k ($k \geq 1$) is a **time-varying confounder** for the effect of A_k on Y if it simultaneously satisfies:

1. L_k is a common cause of A_k and Y (or future treatment–outcome paths).
2. L_k is itself caused by past treatment A_{k-1} (or earlier treatments).

The quintessential example is an intermediate outcome (e.g., CD4 cell count) that is measured repeatedly, is affected by prior antiretroviral therapy, and also predicts both future treatment assignment (physicians adjust the regimen based on CD4) and the ultimate outcome (AIDS or death).

6.1 The Feedback Loop

The defining feature of time-varying confounding is a **treatment–confounder feedback loop**:

$$A_{k-1} \rightarrow L_k \rightarrow A_k \rightarrow Y.$$

Here, L_k lies on the causal path from past treatment to future treatment (through its effect on A_k) and also on an independent path to the outcome. This creates a situation in which L_k is simultaneously:

- A **confounder** for the $A_k \rightarrow Y$ relationship (must be adjusted for), and
- An **intermediate outcome** (mediator) on the $A_{k-1} \rightarrow Y$ path (should not be adjusted for if we want to capture the full effect of past treatment).

This double role is precisely what makes time-varying confounding so challenging. Standard regression-based methods that condition on $\bar{L}_k$ cannot handle it correctly (see Chapter 20 for a detailed analysis). The methods of Chapter 21 — the g-formula, IP weighting, and g-estimation — are specifically designed to handle time-varying confounders in the presence of feedback.

Illustrative example (HIV treatment):

Suppose A_k is antiretroviral therapy at time k , L_k is CD4 count at time k , and Y is death.

- A_{k-1} raises L_k (CD4 increases due to prior treatment).
- L_k influences A_k (physicians prescribe based on CD4).
- L_k predicts Y (low CD4 $\rightarrow$ higher mortality).

Standard regression conditioning on $\bar{L}_k$ partially blocks the causal path $A_{k-1} \rightarrow L_k \rightarrow \dots \rightarrow Y$, creating an under-adjustment for the total effect of $\bar{A}$. At the same time, not conditioning on L_k

leaves open the backdoor path $A_k \leftarrow L_k \rightarrow Y$, creating confounding for the effect of A_k . There is no way to resolve this tension within standard regression.

Why standard methods fail: By conditioning on a time-varying confounder L_k that is affected by past treatment, we introduce collider bias (see Chapter 8) along the path $A_{k-1} \rightarrow L_k \leftarrow U \rightarrow Y$, where U is any unmeasured common cause of L_k and Y . This is explained in detail in Chapter 20.

7 Summary

- The **time-varying treatment** framework models interventions indexed by time k , with treatment history $\bar{A}_k = (A_0, \dots, A_k)$ and covariate history $\bar{L}_k = (L_0, \dots, L_k)$.
- A **treatment strategy** g specifies how treatment decisions are made at each time point, possibly as a function of past history (dynamic strategy) or not (static strategy).
- The counterfactual Y^g is the potential outcome under strategy g ; its mean $E[Y^g]$ is the primary estimand.
- **Sequential exchangeability** ($Y^g \perp\!\!\!\perp A_k \mid \bar{A}_{k-1} = \bar{g}_{k-1}, \bar{L}_k$) and **sequential positivity** are the two key identifying assumptions.
- Not all strategies may be identifiable; the analysis should be restricted to strategies for which both conditions hold.
- **Time-varying confounders** — variables that are caused by past treatment and that cause future treatment and the outcome — create a feedback loop that cannot be handled by standard regression.

8 References

Hernán, Miguel A, and James M Robins. 2020. *Causal Inference: What If*. Chapman & Hall/CRC. <https://miguelhernan.org/whatifbook>.