Section V: Extensions

Today, we will discuss:

- Dynamic Treatment Regimes

For your reference, slides are also included for the following extensions:

- Survival outcomes with censoring
- Multicategory Treatment

Survival outcomes with censoring

Dynamic Treatment Regimes

Dynamic Treatment Regimes (DTRs)

- Motivation : treatment of chronic illness
  - Some examples: HIV/AIDS, cancer, depression, schizophrenia, drug and alcohol addiction, ADHD, etc.
  - Multistage decision making problem
  - Longer-term treatment requires consideration and tradeoff of present versus longer term benefit.

Dynamic Treatment Regimes

- Operationalize multistage decision making via as sequence of decision rules
  - One decision rule for each time (decision) point
  - A decision rule is a function inputs patient history and outputs a recommended treatment
- Aim to optimize some cumulative clinical outcome
  - Survival time
  - Depression test scores
  - Indicator of no myocardial infarction within 30 days ...

Dramatized Example

- Addiction management example inspired by the ExTENd and COMBINE trials (Murphy et al, 2007)

- Devising two-time point treatment strategy for alcohol dependent patients.
  - Initial treatment choices Naltrexone (NTX) and Combined Behavioral Intervention (CBI).
  - At six-months responders classified as responders or non-responders.
  - For responders to initial treatment, followup treatment choices are telephone monitoring (TEL) and telephone monitoring + counseling (TEL+Counseling).
  - For non-responders to initial treatment, followup treatment choices are switch initial treatments (NTX $\leftrightarrow$ CBI), or step-up initial treatment CBI + NTX + Enhanced monitoring (CBI + NTX +EM).

Dramatized Example

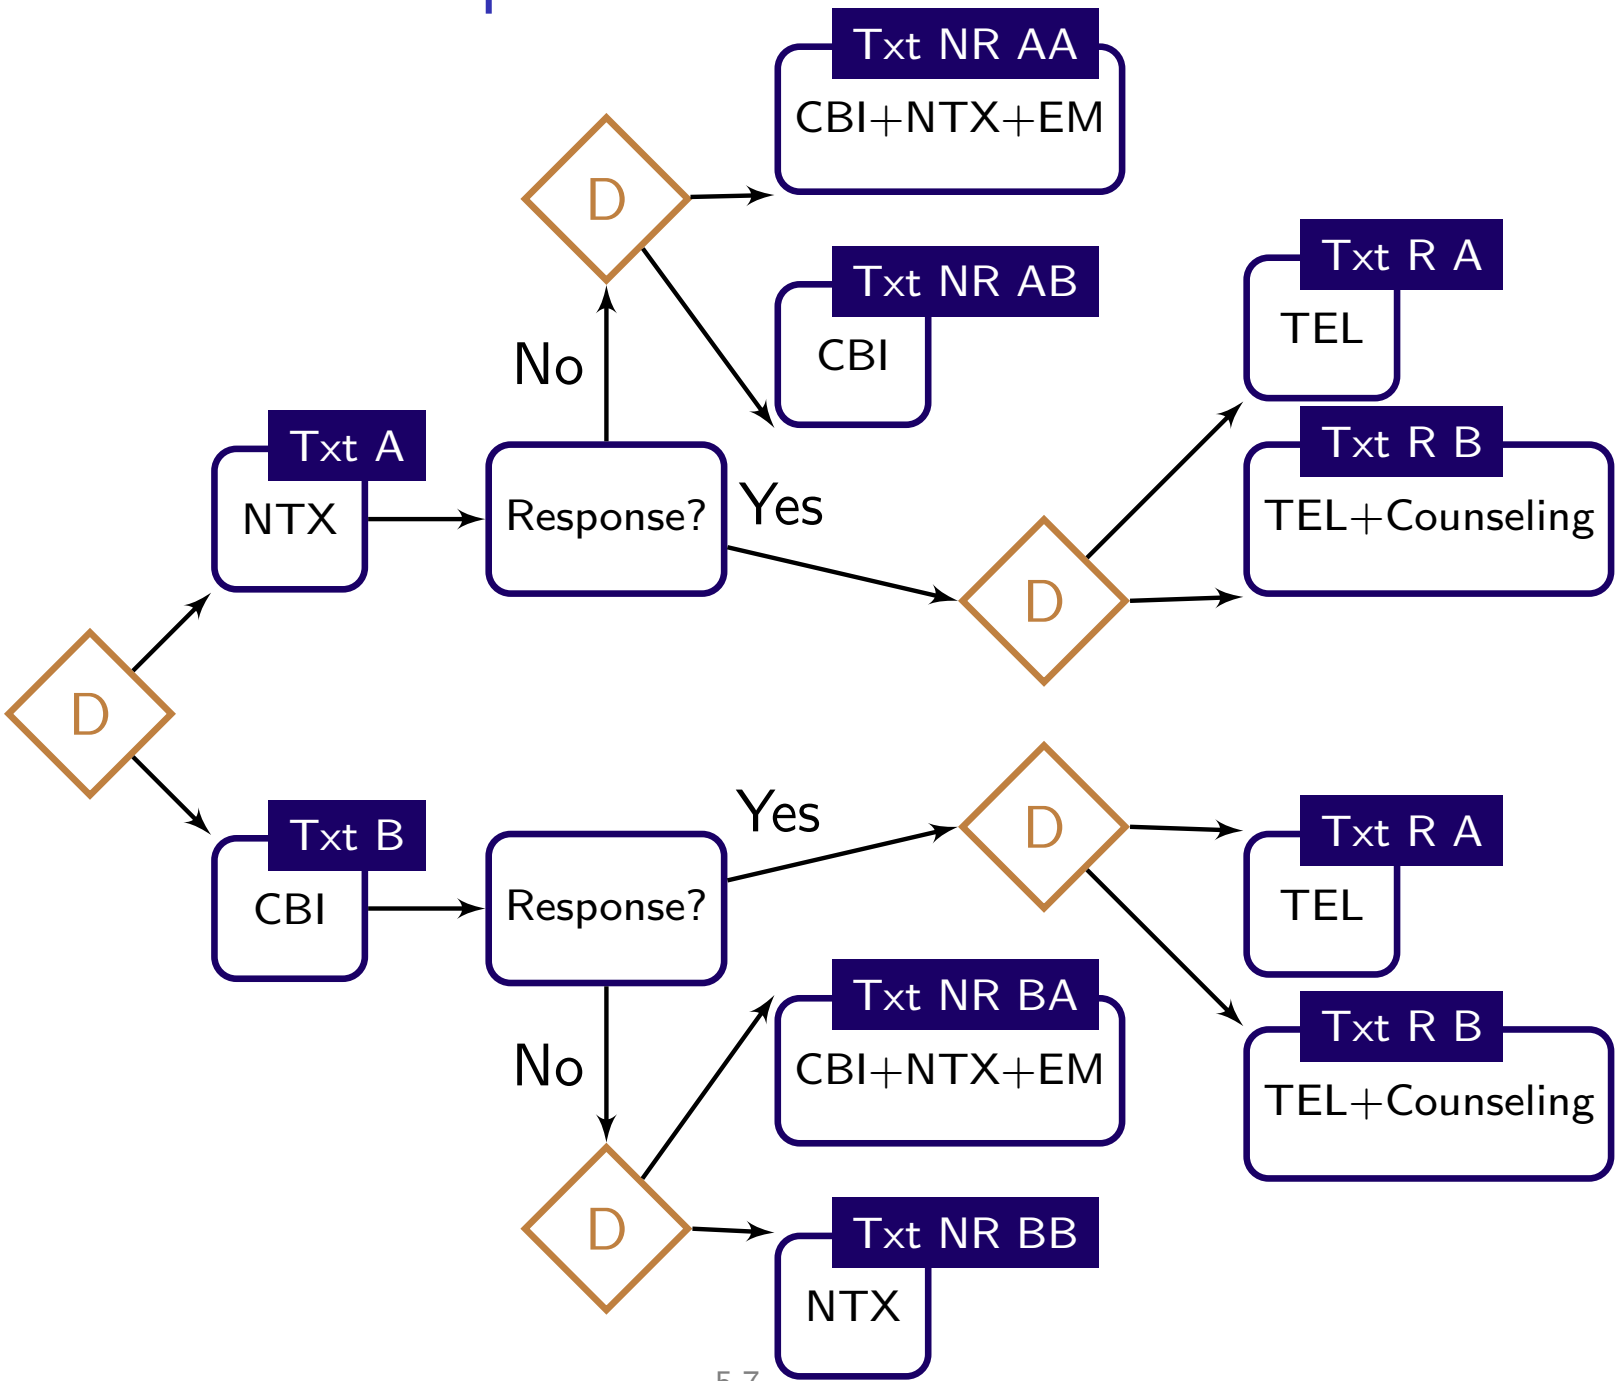

Dramatized Example

- H_j denote history at stage j.
- At presentation: Baseline variables x_1; accrued information $h_1 = x_1$
  - Decision point 1: Two treatment options {NTX, CBI}; rule 1: $d_1(h_1) \Rightarrow d_1 : h_1 \rightarrow \{\text{NTX, CBI}\}$
  - Between decisions 1 and 2: Collect additional information x_2, including responder status
  - Accrued information $h_2 = \{x_1, \text{treatment at decision 1}, x_2\}$
  - Decision point 2: Four options

Optimal Dynamic Treatment Regimes

- Examples of treatment regimes: Prescribe NTX initially; then assign TEL to responders; and assign step-up to non-responders.
- Optimal DTR d^* leads to the lowest expected outcome among all possible regimes

Challenges in Estimating Optimal DTRs: Delayed Effects

- The therapy with the higher proportion of responders might have other effects that render subsequent treatments less effective in regard to the final response.
- The therapy with lower proportion of responders may not appear best initially but may have enhanced long term effectiveness when followed by a particular maintenance treatment.
- Must consider the entire sequence of decisions
- Must accommodate intermediate information including prior treatments into current treatment choice.

Sequential Multiple Assignment Randomized Trial (SMART)

- Due the the aforementioned challenges, it would be ideal to adopt a particular design to best estimate the optimal DTRs
- SMART: designed for estimation of optimal DTRs
- Randomize individuals to the treatment options at each decision point
- Take advantage of sequential randomization to eliminate confounding
- Collect both initial and intermediate information on possible tailoring variables

Murphy (*Stat in Med*, 2005)

SMART Trial

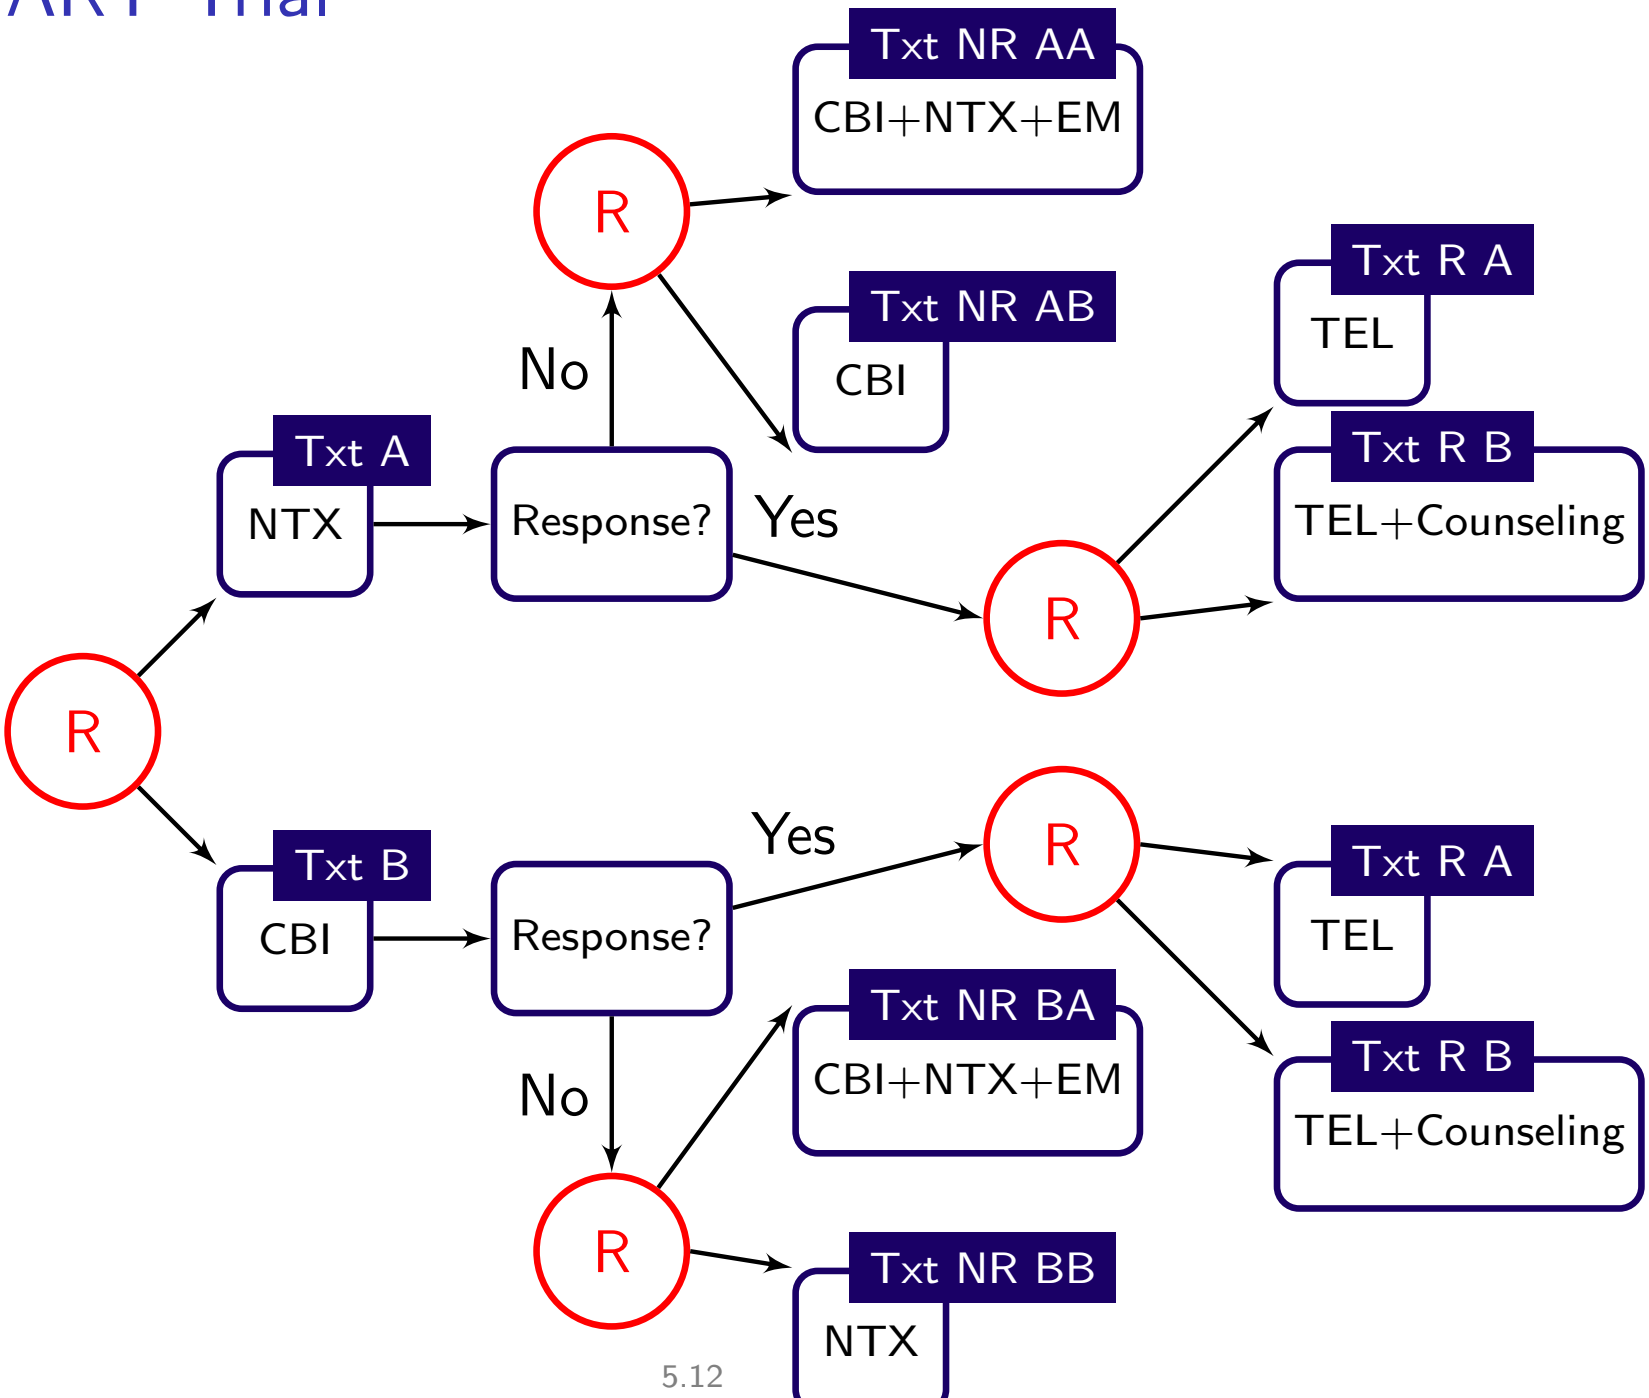
R
Txt A
NTX
Response?
No
R
Txt NR AA
CBI+NTX+EM
Txt NR AB
CBI
Yes
R
Txt R A
TEL
Txt R B
TEL+Counseling
R
Txt B
CBI
Response?
Yes
R
Txt R A
TEL
Txt R B
TEL+Counseling
No
R
Txt NR BA
CBI+NTX+EM
Txt NR BB
NTX

Data

- (X_1, A_1, X_2, A_2, D) for each individual
  X_k: Observations available at stage k
  A_k: Treatment at stage k
  D: Primary outcome
  H_k: History at stage k, $H_1 = X_1,\ H_2 = (X_1, A_1, X_2)$

- The regime, $d = \{d_1, d_2\},\ d_k : \mathcal{H}_k \to \mathcal{A}_k$, should have the lowest $E^d(D)$, the expected outcome if all patients are assigned treatment according to d

Dynamic Programming

- Optimal regime d^* can be derived using dynamic programming (Bellman, 1957)
  - Define
    - $Q_2(h_2, a_2) \triangleq E\left(D \middle| H_2 = h_2, A_2 = a_2\right)$
    - $\tilde{D} \triangleq \min_{a_2} Q_2(H_2, a_2)$
    - $Q_1(h_1, a_1) \triangleq E\left(\tilde{D} \middle| H_1 = h_1, A_1 = a_1\right)$
  - $d_j^*(h_j) = \arg\min_{a_j \in \{0,1\}} Q_j(h_j, a_j)$

Constructing a DTR from Data: Q-learning

- When system dynamics are known dynamic programming yields the optimal DTR, but we only have data
- Q-learning: data-driven analog of dynamic programming: replaces conditional expectations with regression models
- Backwards and recursively estimates the Q-function.
- The estimated optimal sequence of decision rules

$$\hat{d}_j(h_j) = \underset{a_j \in \{0,1\}}{\operatorname{argmin}} \hat{Q}_j(h_j, a_j).$$

- An extension of regression to sequential treatments.

Summary

- An extremely active area of research
- Data from SMART designs can be used to construct optimal DTRs
- Q learning is a common method, though it has some drawbacks, e.g., require correct specified models
- Many other methods have been developed.

Survival Outcomes with Censoring

- Interested in time-to-event outcome.
- Observe independently and identically distributed training data $(X_i, A_i, D_i, \Omega_i), i = 1, \ldots, n$.

  X: baseline variables, $X \in \mathbb{R}^p$,
  A: binary treatment options, $A \in \{0, 1\}$,
  D: observed event time.
  Ω: censoring indicator $\Omega_i = I(T_i \leq C_i)$.

- $D = \min(T, C)$: T survival time, C censoring time.
- Randomized study with known randomization probability of the treatment.

Survival Outcomes with Censoring

Survival Outcomes with Censoring

- Two possible objectives
  - Maximize the probability of surviving beyond a landmark time;
  - Maximize restricted mean survival time.

Probability of surviving beyond a landmark time

Let T be the event time. Let $D = I(T < t_0)$ be an indicator that the event occurs before a landmark time t_0.

- Estimate $E(D|A, X)$ using a regression method suitable for time-to-event outcomes (e.g. Cox regression with treatment-by-covariate interactions). This may need to be paired with a baseline hazard estimate.
- Consider performing analyses for different choices of t_0; typically X more weakly predicts treatment effect for larger t_0.

Cox (*JRSSB*, 1972)

Restricted mean survival time

- Regression modeling approach: inverse probability of censoring weighted (IPW) Q-learning:
  - $E(D|A,X)$ is modeled using treatment-by-covariate interactions, accounting for the probability of being censored.
- Outcome weighted learning approach:
  - Replace D_i by $\Omega_i D_i / \hat{S}_C(D_i|A_i,X_i)$ in the outcome weighted learning for uncensored data, where $\widehat{S}_C(D|A,X)$ is the estimated conditional survival function of C given (A,X).
  - Doubly robust idea: identify a double robust version of the value function using the augmented IPW estimators.

Goldberg & Kosorok (*Annals of Stat.*, 2012); Zhao et al (*Biometrika*, 2015)

Evaluation in the censoring data setup

Estimate performance measures empirically using inverse-probability-of-censoring weights. (Model-based estimates require no modification.)

Multicategory Treatment

Multicategory Treatment

- Multiple treatments of interest, $A = 0, 1, \ldots, K$, e.g., $K = 2$ in depression data
- $d^*(x) = \text{argmin}_{k=0,\ldots,K}\, \mu(k, x)$.
- Posit a regression model

$$E(D|A, X) = \mu(A, X; \beta)$$

  and estimate $\hat{\beta}$.

- The estimator for the optimal treatment regime

$$\hat{d}_n(x) = \underset{k=0,\ldots,K}{\text{argmin}}\, \mu(k, x; \hat{\beta}_n).$$

- Other methods under development.