

# 67. Recent Innovations in Dynamic Treatment Regimes

---

- 1:45 - 2:15 *An Introduction to Dynamic Treatment Regimes* - Marie Davidian
- 2:15 - 2:45 *Estimating Mean Response of Treatment Duration Regimes in an Observational Study* - Butch Tsiatis
- 2:45 - 3:15 *Hypothesis Testing and Dynamic Treatment Regimes* - Susan Murphy

# An Introduction to Dynamic Treatment Regimes

---

Marie Davidian  
Department of Statistics  
North Carolina State University



<http://www.stat.ncsu.edu/~davidian>

# Outline

---

1. What is a dynamic treatment regime?
2. How can we make inferences on dynamic treatment regimes?
3. Dynamic treatment regimes in oncology
4. Constructing dynamic treatment regimes
5. Concluding remarks

# 1. What is a dynamic treatment regime?

---

**Clinical practice:** Treatment of *ongoing illness* is itself *ongoing*

- Providers *do not* make a *once-and-for-all* decision, e.g., “take drug A no matter what happens to you, just keep taking it”
- Rather, providers *manage* a patient’s illness
- Providers routinely *adjust*, *change*, *add*, or *discontinue* treatment based on *progress*, *side effects*, *patient burden*, *compliance*, etc.
- Providers think of this as “*individualizing*” treatment to the patient

**That is:** Treatment in practice involves *decisions made sequentially over time* based on *accruing observations on the patient*

- Suggests *thinking about* and *studying* treatment this way . . .

# 1. What is a dynamic treatment regime?

---

## How are these decisions made?

- *Clinical judgment* based on experience
- *Patient preference*
- *Practice guidelines* based on pieced-together *clinical evidence* and *expert opinion*, e.g., NIH guidelines for treatment of HIV-infection at <http://www.aidsinfo.nih.gov/>

## Dynamic treatment regime: Aka *adaptive treatment strategy*

- A set of *sequential decision rules*, each of which dictates how to make the decision on what to do next for a patient based on *observation* of the patient up to that point
- Allows treatment to be “*individualized*” to the patient through a systematic set of rules that *operationalize* clinical practice

# 1. What is a dynamic treatment regime?

---

**Dynamic treatment regime:** More precisely

- A sequence of *decision points* at which *decisions on treatment* are made
- At each point, the *next step* of treatment is determined according to *information up to that point* on the patient ...
- ...based on a *decision rule* that takes these *variables* as *input* and *outputs* the *next treatment step* for the patient
- May be thought of as an *algorithm* that dictates how treatment of a patient should proceed over time

**A concrete example ...**

# 1. What is a dynamic treatment regime?

---

## Goals of cancer therapy:

- *Induce* remission of disease, usually using powerful chemotherapeutic agents
- *If remission*, then *maintain* remission as long as possible before relapse/recurrence, e.g., by administering additional agents that *intensify or maintain* the effects of the initial induction therapy
- *If no remission*, then maybe try *something else* to induce remission

**Primary outcome of interest:** A time-to-event, e.g., *disease-free survival time*

**Note:** We are currently discussing *dynamic treatment regimes* in the context of a *single patient* (we will talk about *studies* involving many subjects shortly. . . )

# 1. What is a dynamic treatment regime?


---

**A particular dynamic treatment regime:** For a *given patient*

- *Step 1*: Treat with one or more courses of first-line *induction* chemotherapy  $A$
- *Intermediate outcome*: Observe whether “*response*” occurs
- “*Response*” may be defined as complete or partial remission, degree of tumor shrinkage, etc.
- *Step 2*: If “*response*” occurs, give *maintenance* therapy  $B$  ...
- ... else, if “*response*” *does not* occur (so  $A$  *did not induce* a response), try *second-line* therapy  $B'$

**Decision rule:** The *decision rule* to determine the step 2 treatment takes the variable “*response or not?*” as *input*

**Schematically:** The specific regime “Give first-line induction therapy  $A$  followed by maintenance  $B$  if response else if no response give second-line therapy  $B'$ ”



*Step 1*

*Step 2*

(Induction Trt)

(Intermediate Outcome)

(Maintenance or Second-line Trt)

# 1. What is a dynamic treatment regime?

---

**Important:** Individuals following *the same regime* can have *different realized treatment experiences*

- *Subject 1*: Receives  $A$ , responds, receives  $B$
- *Subject 2*: Receives  $A$ , does not respond, receives  $B'$
- *Both* subjects' experiences are *consistent with* following this regime

**Important, part 2:** *Do not confuse* the *regime* with the *possible realized experiences* that can result from following it

- There *are NOT two regimes*! I.e., “ $A$  followed by response followed by  $B$ ” and “ $A$  followed by no response followed by  $B'$ ” are *not regimes*. They are possible *results* of following the *single regime*!
- The *regime* is the *algorithm* that dictates how to treat a patient *over time*

## 2. Inferences on dynamic treatment regimes

---

**Options:** There may be *more than one* possible regime

- More than one possible *first-line induction* treatment (*Step 1*), e.g., two options  $A_1$  and  $A_2$
- More than one possible *maintenance* treatment if response occurs (*Step 2*), e.g., two options  $B_1$  and  $B_2$
- More than one possible *second-line induction* treatment if no response occurs (*Step 2*), e.g., two options  $B'_1$  and  $B'_2$
- *In general*: The number and types of options at each step *need not even be the same*

## 2. Inferences on dynamic treatment regimes

---

### Eight possible regimes:

1.  $A_1$  followed by  $B_1$  if response, else  $B'_1$
2.  $A_1$  followed by  $B_1$  if response, else  $B'_2$
3.  $A_1$  followed by  $B_2$  if response, else  $B'_1$
4.  $A_1$  followed by  $B_2$  if response, else  $B'_2$
5.  $A_2$  followed by  $B_1$  if response, else  $B'_1$
6.  $A_2$  followed by  $B_2$  if response, else  $B'_2$
7.  $A_2$  followed by  $B_1$  if response, else  $B'_1$
8.  $A_2$  followed by  $B_2$  if response, else  $B'_2$

### Natural questions:

- What would be the *mean outcome* (e.g., *mean survival time*) if the *population* were to *follow* a particular regime?
- How do these mean outcomes *compare* among the possible regimes?

### How might we address such issues?

## 2. Inferences on dynamic treatment regimes

---

### Can't we learn about this based on a series of previous trials?

- In one trial,  $A_1$  was compared against  $A_2$  in terms of *response rate*
- In another trial,  $B_1$  and  $B_2$  were compared on the basis of *survival* in subjects who *responded* to their first-line chemotherapy
- In yet another,  $B'_1$  and  $B'_2$  were compared (*survival*) in subjects for whom first-line therapy *did not induce response*
- Can't we just "*piece together*" the results from these separate trials to figure out the "*best regime*?"
- E.g., figure out the best "*A*" treatment for inducing response and then the best "*B*" and "*B'*" treatments for prolonging survival?
- Wouldn't the regime that uses these have to have the "*best*" mean outcome?

## 2. Inferences on dynamic treatment regimes

---

### A problem with this: *Delayed effects*

- E.g.,  $A_1$  may yield *higher proportion of responders* than  $A_2$  but may also have other effects that render subsequent intensification treatments ( $B$ ) *less effective* in regard to *survival*
- $\implies$  Must study *entire regimes*

### So how can we do this?

- Design a *clinical trial* in which subjects are *randomized* to follow different regimes – we will focus on this *next*...
- Use *observational* follow-up data (*somehow*), e.g., from a registry or other database, where *treatments actually received* over time have been recorded (with other information) for each subject – *Butch* will discuss a case study

## 2. Inferences on dynamic treatment regimes

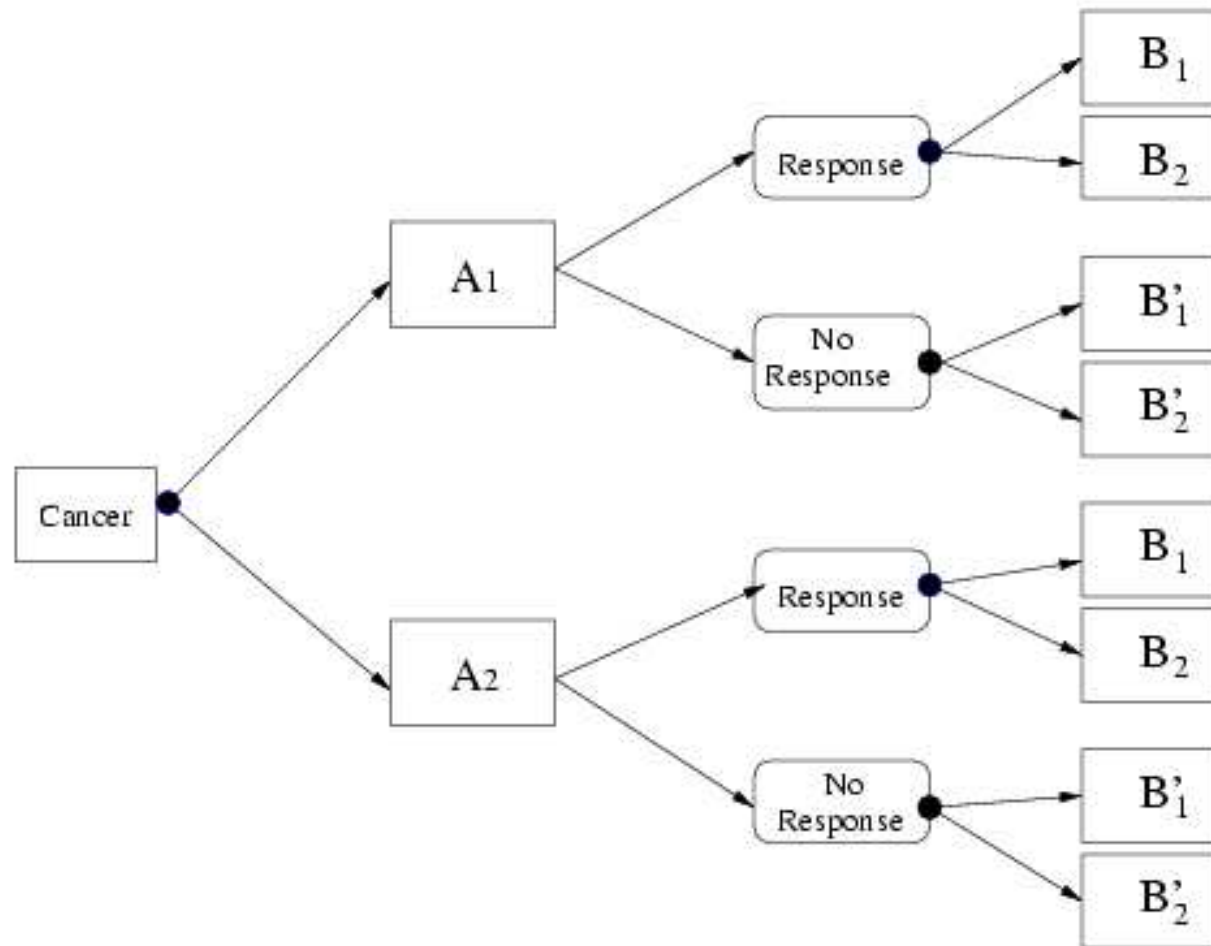
---

### Clinical trials to study dynamic treatment regimes:

- An *eight-arm* trial – subjects randomized to the  $j$ th arm *follow* the  $j$ th regime
- A *sequentially-randomized* trial (*next slide...*)
- How to *analyze* the *outcome data* to compare regimes in such trials?  
What *else* can be *learned* from such trials?

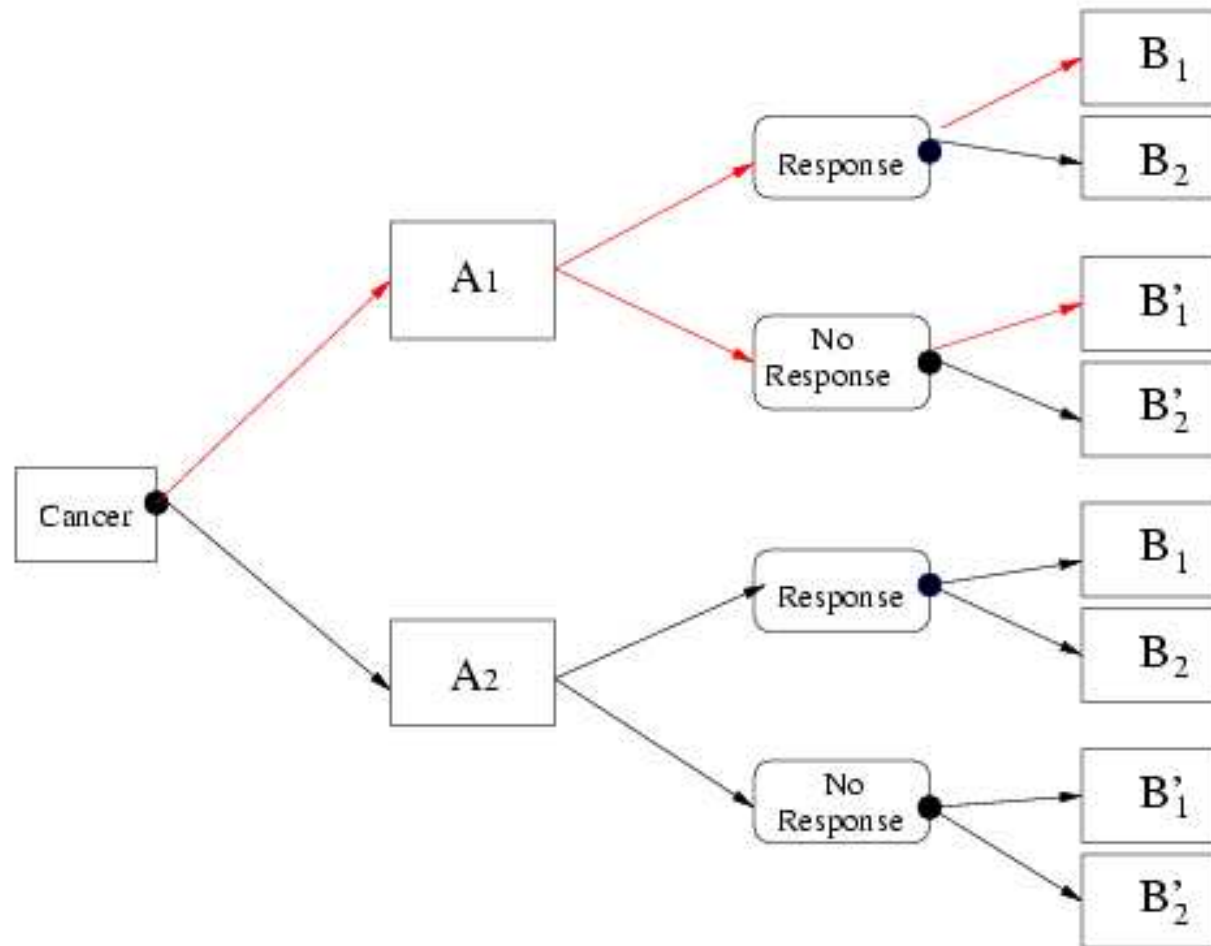
## 2. Inferences on dynamic treatment regimes

“**SMART:**” Sequential Multiple Assignment Randomized Trial, e.g., Lavori and Dawson (2003) (Randomization at ●s)



## 2. Inferences on dynamic treatment regimes

**In red:** Regime “ $A_1$  followed by  $B_1$  if response else  $B'_1$ ”



## 2. Inferences on dynamic treatment regimes

---

### Remarks:

- *Susan* has developed a *general framework* for designing SMARTs:  
Murphy, SA. (2005). An experimental design for the development of adaptive treatment strategies, *Statistics in Medicine* **24**, 1455-1481.
- One can determine the *randomization probabilities* at each step so that the numbers of subjects ending up with *treatment experiences consistent with each regime* is the *same*, as expected in the *eight-arm* trial with “*up-front*” *randomization to the regimes* with equal probabilities
- *Thus*, there is no *conceptual difference* between randomizing *up-front* or *sequentially*
- *However* there are special considerations for *analysis*...

## 2. Inferences on dynamic treatment regimes

---

### Estimation of mean outcome (e.g., mean survival):

- Usual approach under *up-front randomization*: estimate mean for regime  $j$  by *sample average outcome* based on subjects randomized to regime  $j$  only
- *However*: Subjects will have *realized experiences consistent with more than one regime!*
- E.g., *Realized treatment experience*

$$A_1 \Rightarrow \text{Response} \Rightarrow B_1$$

is *consistent with BOTH* regimes

- $A_1$  followed by  $B_1$  if response, else  $B'_1$
  - $A_1$  followed by  $B_1$  if response, else  $B'_2$
- This can be *exploited* to improve precision...

## 2. Inferences on dynamic treatment regimes

---

### Demonstration:

- A certain kind of SMART is common in *oncology*...
- ...but way these trials are usually analyzed *does not* focus on comparing the embedded *dynamic treatment regimes*
- Such an analysis is proposed in

Lunceford JK, Davidian M, Tsiatis AA. (2002). Estimation of survival distributions of treatment policies in two-stage randomization designs in clinical trials. *Biometrics* **58**, 48–57.

and demonstrates the *general principle* of how to exploit realized experiences consistent with more than one regime...

### 3. Dynamic regimes in oncology

---

#### Cancer and Leukemia Group B (CALGB) Protocol 8923:

Double-blind, placebo-controlled trial of 338 elderly subjects with acute myelogenous leukemia (AML) with *two randomizations*

- Subjects randomized to either *standard induction chemotherapy*  $A_1$  OR *standard induction therapy + granulocyte-macrophage colony-stimulating factor (GM-CSF)*  $A_2$  (*Step 1* options)
- *If response*, subjects randomized to  $B_1, B_2 =$  *intensification* treatments I, II *Step 2* options
- *If no response*, only *one Step 2* option: follow-up with physician
- All subjects followed for the *outcome survival time*

# 3. Dynamic regimes in oncology


---

## Four possible regimes:

1.  $A_1$  followed by  $B_1$  if response, else follow up with physician =  $A_1B_1$
2.  $A_1$  followed by  $B_2$  if response, else follow up with physician =  $A_1B_2$
3.  $A_2$  followed by  $B_1$  if response, else follow up with physician =  $A_2B_1$
4.  $A_2$  followed by  $B_2$  if response, else follow up with physician =  $A_2B_2$

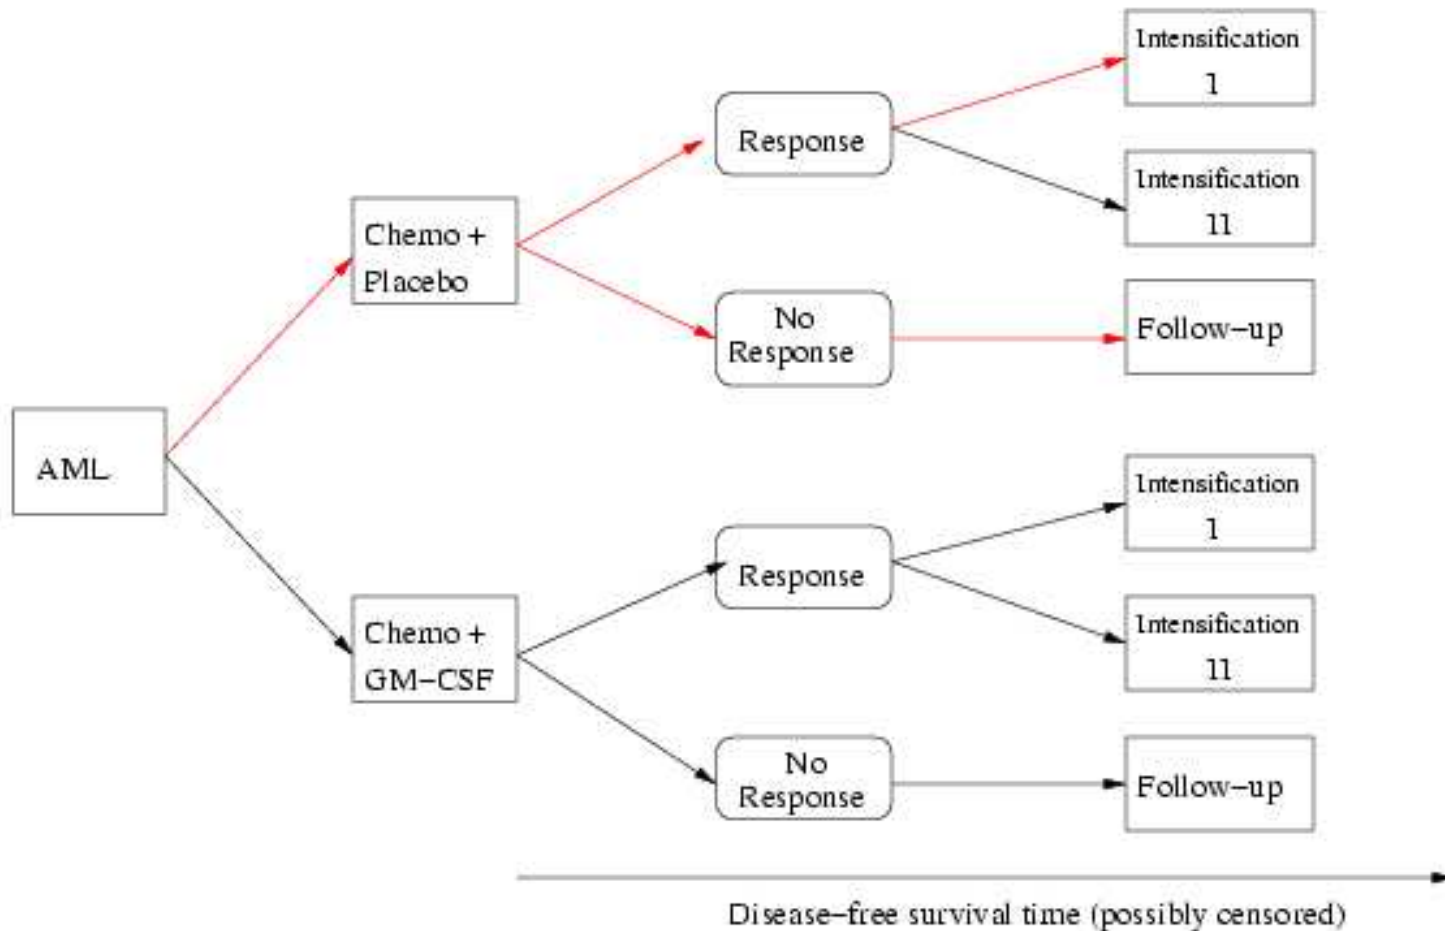
# 3. Dynamic regimes in oncology

**Schematic of CALGB 8923:** Randomization at ●s



### 3. Dynamic regimes in oncology

Regime  $A_1B_1$ :



# 3. Dynamic regimes in oncology

---

## Standard analysis:

- Compare *response rates* to  $A_1$  and  $A_2$
- Compare *survival* between  $B_1$  and  $B_2$  among *responders*
- Compare *survival* between  $A_1$  and  $A_2$ , regardless of subsequent response/randomization

**Problem:** Does not address *directly* comparison of the *dynamic treatment regimes*

**Demonstration:** For each regime  $A_j B_k$ ,  $j = 1, 2$ ,  $k = 1, 2$

- *Estimate* the *mean survival time* under regime  $A_j B_k$
- I.e., estimate mean survival if the *entire AML population were to follow* regime  $A_j B_k$

### 3. Dynamic regimes in oncology

---

**Basic idea:** To *estimate the mean* for  $A_j B_k$ , use data from all subjects whose *realized experience* is *consistent with* having followed  $A_j B_k$

- $A_j \Rightarrow$  *response*  $\Rightarrow B_k$
- $A_j \Rightarrow$  *no response*  $\Rightarrow$  follow up with physician
- *Combine* survival times from these subjects in an *appropriate way*...

**Consider  $A_1$  only ( $A_2$  analogous):** Suppose *responders* are randomized to  $B_1$  or  $B_2$  with probability  $1/2$

- *Nonresponders* to  $A_1 \Rightarrow$  follow up
- *Half of responders* get  $B_1$ , *half* get  $B_2$
- The half who end up randomized to  $B_2$  have *missing* survival times as far as  $A_1 B_1$  is concerned

### 3. Dynamic regimes in oncology

---

**Result:** To estimate *mean survival* for  $A_1B_1$  from the trial

- Use a *weighted average*
- The *nonresponders* represent themselves  $\Rightarrow$  weight = 1
- Each *responder* represents him/herself and another similar subject who got randomized to  $B_2 \Rightarrow$  weight = 2
- Usual “*inverse probability weighting*” for missing data
- To estimate mean for  $A_1B_2$ , switch the roles
- *Note*: Survival times from *nonresponders* are used to estimate the means for *both*  $A_1B_1$  *and*  $A_1B_2$

**In general:** This idea can be *extended* to any number of steps and numbers of options at each step

### 3. Dynamic regimes in oncology

---

**In symbols:** Suppose  $n$  subjects end up randomized to  $A_1$

$T_i$  = survival time for subject  $i$ ,  $i = 1, \dots, n$ ,

$R_i = 1$  if  $i$  responds to  $A_1$ ,  $R_i = 0$  if not

$Z_i = 1$  for a responder randomized to  $B_1$ ,  $Z_i = 2$  for  $B_2$

$P(Z_i = 1 | R_i = 1) = \pi$  (= 1/2 in previous)

**Estimators:**  $n^{-1} \sum_{i=1}^n Q_i T_i$  or  $\left( \sum_{i=1}^n Q_i \right)^{-1} \sum_{i=1}^n Q_i T_i$ ,

$$Q_i = 1 - R_i + R_i I(Z_i = 1) \pi^{-1}$$

- $Q_i = 0$  if  $i$  is inconsistent with  $A_1 B_1$  (i.e, is consistent with  $A_1 B_2$ )
- $Q_i = 1$  if  $R_i = 0$
- $Q_i = \pi^{-1}$  if  $R_i = 1$  and  $Z_i = 1$
- To estimate  $S(t) = P(T_i > t)$ , estimate  $F(t) = 1 - S(t)$  by replacing  $T_i$  by  $I(T_i \leq t)$

# 3. Dynamic regimes in oncology

---

## Remarks:

- Subjects may *die* before having a chance to respond – *nonresponders* at the time of death ( $R_i = 0$ )
- Survival time may be *right-censored* – can incorporate *inverse probability of censoring* weighting
- *Consent of responders*: In CALGB 8923, 10% subjects who *did respond refused to be randomized* at the second stage
- “*Intention to treat*” perspective: Consider instead *offering*  $A_j$  followed by *offering*  $B_k$  if response else follow up
- *Redefine*, e.g., “ $A_1$  followed by  $B_k$  if *response and consent* ( $R_i = 1$ ) else ( $R_i = 0$ ) follow up” (so compare without regard to differential consent rates)
- ... As opposed to attempting to ask the original *causal* question, with this *noncompliance* as a nuisance ( $\Rightarrow$  *observational study*)

# 3. Dynamic regimes in oncology

---

## Issues in designing and analyzing SMARTs:

- *Sample size calculation*
- *Explosion* as number of steps *increases*?
- One possibility: Exploit *experimental design* principles, e.g., *fractional factorial*?

## 4. Constructing dynamic treatment regimes

---

**A bigger question:** How can we *construct* better dynamic regimes?

- The foregoing applies to studying and comparing *simple regimes* that are *preconceived*
- Can results of such trials be used to develop *more refined algorithms* that take *additional evolving information* into account in the *rules* to better “*individualize*” treatment?
- What are the *options* at each step?
- What should be the *timing* of the steps?
- What *information* (*variables*) should be used to select an option at and/or timing for each step?
- How should many such *variables* be *summarized* to create *decision rules* that make sense to clinicians and patients?
- What is the “*best*” sequence of steps (i.e., the *optimal regime*)?

# 4. Constructing dynamic treatment regimes

---

## SMARTs:

- Can form the basis for *comparing* specific simple regimes. . .
- . . . but also *have potential* in *constructing* regimes
- *Key: Randomization* at each step provides a basis for identifying important “*tailoring variables*” at each step; i.e., variables that *affect outcome differentially* by treatment at that step. . .
- . . . and hence should be incorporated in the *decision rule* at that step
- Methods for identifying *tailoring variables* must incorporate *effect of future treatment decisions* when evaluating *present treatment decision* need to be *adapted* to this setting
- *Susan* will discuss methods useful for constructing of dynamic treatment regimes on the basis of data from (sequentially) randomized trials

## 5. Concluding remarks

---

- *Dynamic treatment regimes operationalize* how clinicians *practice* medicine
- *Statisticians* should be open to thinking about treatment as *sequential multi-stage decision-making*
- We should be encouraging our clinician colleagues to think similarly and to conduct studies to evaluate *regimes* rather than *single steps* of treatment
- Methods for doing this *exist*!
- We should be interested in *developing and evaluating* methods for *constructing better dynamic regimes*

**These slides at:** <http://www.stat.ncsu.edu/~davidian>

# 6. Advertisement

---

**Interested in learning more?** You're *in luck!*

*SAMSI 2007 Summer Program on  
Challenges in Dynamic Treatment Regimes and Multistage Decision-Making*

- *Program leaders*: Susan Murphy, Butch Tsiatis, Marie Davidian, Dan Scharfstein, Joelle Pineau
- Two-week intensive study of all of this, *June 18–29, 2007*
- Statisticians, applied mathematicians, computer scientists, ORers, engineers, . . .
- Tutorials, workshops, working groups

<http://www.samsi.info/programs/2007adaptivetreatmentprogram.shtml>



# Some references

---

## SMART analysis stuff:

Lunceford JK, Davidian M, Tsiatis AA. (2002). Estimation of survival distributions of treatment policies in two-stage randomization designs in clinical trials. *Biometrics* **58**, 48–57.

Murphy, SA. (2005). An experimental design for the development of adaptive treatment strategies, *Statistics in Medicine* **24**, 1455-1481. See also <http://www.stat.lsa.umich.edu/~samurphy/research.html>

Thall PF, Millikan RE, Sung HG. (2000). Evaluating multiple treatment courses in clinical trials. *Statistics in Medicine* **30**, 1011-1128.

Wahed AS, Tsiatis AA. (2004). Optimal estimator for the survival distribution and related quantities for treatment policies in two-stage randomization designs in clinical trials. *Biometrics* **60** 124-133.

Wahed AS, Tsiatis AA. (2006). Semiparametric efficient estimation of survival distribution for treatment policies in two-stage randomization designs in clinical trials with censored data. *Biometrika* **93**, 163-177.

# Some references

---

## Constructing regimes stuff:

Murphy SA, Oslin DW, Rush JA, Zhu J for MCATS. (2006).

Methodological challenges in constructing effective treatment sequences for chronic psychiatric disorders.

*Neuropsychopharmacology* **32**, 257–262 .

Murphy SA (2003). Optimal dynamic treatment regimes (with discussion). *JRSS-B* **65**, 331-366.

Murphy SA, van der Laan MJ, Robins JM. (2001). Marginal mean models for dynamic regimes. *JASA* **96**, 1410-1423.

Robins JM (2004). Optimal structural nested models for optimal sequential decisions. In DY Lin, PJ Heagerty (eds.) *Proceedings of the Second Seattle Symposium on Biostatistics*. New York: Springer, 189-326.

# Appendix

---

**One way to formalize:** What are we *estimating*?

- Suppose  $i$  has *potential outcomes* (aka *counterfactuals*)  $T_{11i}, T_{12i}$
- $T_{1ki}$  = survival time  $i$  *would have* if  $i$  *were to follow*  $A_1B_k, k = 1, 2$

**Question of interest:** Estimate *mean survival* if the *entire AML population were to follow* regime  $A_1B_k$

- Distributions of the  $T_{1k}$  represent survival in the population *if all subjects* followed  $A_1B_k, k = 1, 2,$   
 $\Rightarrow$  Want to estimate  $\mu_{1k} = E(T_{1ki})$
- *Similarly*, if interested in the *survival distribution* if all subjects followed  $A_1B_k,$   
 $\Rightarrow$  Want to estimate  $S(t) = P(T_{1ki} > t) = E\{I(T_{1ki} > t)\}$

# Appendix

---

**Of course:** *Do not* observe *both of*  $T_{11i}$ ,  $T_{12i}$  for each  $i$

**Do observe:**  $(R_i, R_i Z_i, T_i)$ ,  $i = 1, \dots, n$ , iid

$T_i$  = survival time for subject  $i$

$R_i = 1$  if  $i$  responds to  $A_1$ ,  $R_i = 0$  if not

$Z_i = k$  for a responder randomized to  $B_k$ ,  $k = 1, 2$ , where

$$P(Z_i = 1 | R_i = 1) = \pi, P(Z_i = 2 | R_i = 1) = 1 - \pi$$

**Consider  $k = 1$ :** Want to estimate  $\mu_{11} = E(T_{11i})$ ,  $k = 1, 2$ , based on *observed data*  $(R_i, R_i Z_i, T_i)$ ,  $i = 1, \dots, n$

- Need to make a *connection* between the *observed data* and the *potential outcomes* ...

- ... to show that  $n^{-1} \sum_{i=1}^n Q_i T_i$  is a *consistent estimator* for  $\mu_{11}$

# Appendix

---

**Connection:** For subjects *randomized* to  $A_1$

- *Assume* that when  $R_i = 0$ ,  $T_{11i}$  and  $T_{12i}$  are *the same*
- Then  $T_i = (1 - R_i)T_{11i} + R_i I(Z_i = 1)T_{11i} + R_i I(Z_i = 2)T_{12i}$

**Want to show:**  $E(Q_i T_i) = E(T_{11i})$ ,  $Q_i = 1 - R_i + R_i I(Z_i = 1) \pi^{-1}$

- *Using*  $R_i(1 - R_i) = 0$ ,  $I(Z_i = 1)I(Z_i = 2) = 0$ , etc.

$$\begin{aligned} E(Q_i T_i) &= E[T_{11i} \{ (1 - R_i) + R_i I(Z_i = 1) \pi^{-1} \}] \\ &= E[T_{11i} E\{ (1 - R_i) + R_i I(Z_i = 1) \pi^{-1} | R_i, T_{11i} \}] \end{aligned}$$

- So *equivalently* want to show

$$E\{ (1 - R_i) + R_i I(Z_i = 1) \pi^{-1} | R_i, T_{11i} \} = 1$$

# Appendix

---

$$\begin{aligned} & E\{(1 - R_i) + R_i I(Z_i = 1)\pi^{-1} | R_i, T_{11i}\} \\ &= E\{(1 - R_i) + R_i I(Z_i = 1)\pi^{-1} | R_i = 0, T_{11i}\}P(R_i = 0 | T_{11i}) \\ &+ E\{(1 - R_i) + R_i I(Z_i = 1)\pi^{-1} | R_i = 1, T_{11i}\}P(R_i = 1 | T_{11i}) \\ &= P(R_i = 0 | T_{11i}) + E\{I(Z_i = 1) | R_i = 1, T_{11i}\}\pi^{-1}P(R_i = 1 | T_{11i}) \\ &= P(R_i = 0 | T_{11i}) + P(R_i = 1 | T_{11i}) = 1 \end{aligned}$$

**Because:** By *randomization*,

$$E\{I(Z_i = 1) | R_i = 1, T_{11i}\} = P(Z = 1 | R = 1, T_{11i}) = P(Z = 1 | R = 1) = \pi$$

$\Rightarrow$  *randomization* ensures  $i$ 's assignment to  $B_1$  *does not depend on*  $i$ 's prognosis

**For  $k = 2$ :** Same argument, now  $Q_i = 1 - R_i + R_i I(Z_i = 2)(1 - \pi)^{-1}$ ,