

I. Dynamic Treatment Regimes in Public Health

9:00-9:05 am	Introduction to Session
9:05-9:35 am	Estimation of Survival Distributions for Treatment Regimes in Two Stage Oncology Trials <i>Marie Davidian</i> , NC State University
9:35-9:40 am	Discussion
9:40-10:25 am	Estimating Mean Response as a Function of Treatment Duration in an Observational Study <i>Anastasios A. (Butch) Tsiatis</i> , NC State University
10:25-10:30 am	Discussion
10:30-10:45 am	<i>Break</i>
10:45-11:45 am	SMART Designs for Developing Dynamic Treatment Regimes <i>Susan A. Murphy</i> , University of Michigan
11:45-noon	Discussion

I. Dynamic Treatment Regimes in Public Health

Objectives of this session:

- Introduce the notion of a *dynamic treatment regime* (or *adaptive treatment strategy*) through two case studies (*Marie*, *Butch*)
- Describe methods for making inference about particular dynamic treatment regimes from *randomized studies* and from *observational data* (*Marie*, *Butch*)
- Describe a general framework for thinking about and *designing* dynamic treatment regimes and in particular for identifying the “*best*” dynamic treatment regime (*Susan*)

I. Dynamic Treatment Regimes in Public Health

Dynamic treatment regime:

- “Individually-tailored” *sequence* of treatment steps
- The *next step* of treatment is determined according to *subject outcomes and information* up to that point
- Consistent with *clinical practice*

I. Dynamic Treatment Regimes in Public Health

Dynamic treatment regime:


- “Individually-tailored” *sequence* of treatment steps
- The *next step* of treatment is determined according to *subject outcomes and information* up to that point
- Consistent with *clinical practice*

Issues:

- What are the *options* at each step?
- What *information* should be used to select an option at each step?
- What should be the *timing* of the steps?
- What is the “*best*” sequence of treatment steps?
- From what kinds of *studies* can we learn about all of this?

Estimation of Survival Distributions for Treatment Regimes in Two Stage Oncology Trials

Marie Davidian
Department of Statistics
North Carolina State University



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

(Joint work with A.A. Tsiatis, J. Lunceford, A. Wahed)

Outline

1. Dynamic treatment regimes for cancer
2. Randomized oncology trials to compare dynamic regimes
3. Case study: CALGB 8923
4. Analysis
5. Wrinkles
6. Discussion
7. Demonstration using potential outcomes
8. References

1. Dynamic treatment regimes for cancer

Goals of cancer therapy:

- *Induce* remission of disease, usually using powerful chemotherapeutic agents
- *Maintain* remission as long as possible before relapse/recurrence, e.g., by administering additional agents that intensify or augment the effects of the initial induction therapy


1. Dynamic treatment regimes for cancer

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
- *Step 2*: If “*response*” occurs, give *maintenance* therapy B . . .
- . . . else, if “*response*” does not occur (so A *did not induce* a response), do something else, e.g., try a *second-line* therapy B'
- “*Response*” typically defined as complete or partial remission, tumor shrinkage, etc.

Primary outcome of interest: E.g., in cancer, *disease-free survival time*

Schematically: The specific regime “Give first-line induction therapy A followed by maintenance B if response else give second-line therapy B' ”



Step 1

Step 2

(Primary Trt)

(Intermediate
Outcome)

(Secondary Trt)

1. Dynamic treatment regimes for cancer

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

1. Dynamic treatment regimes for cancer

Eight possible regimes or strategies:

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

Question: How do these eight regimes *compare* on the basis of *disease-free survival time*?

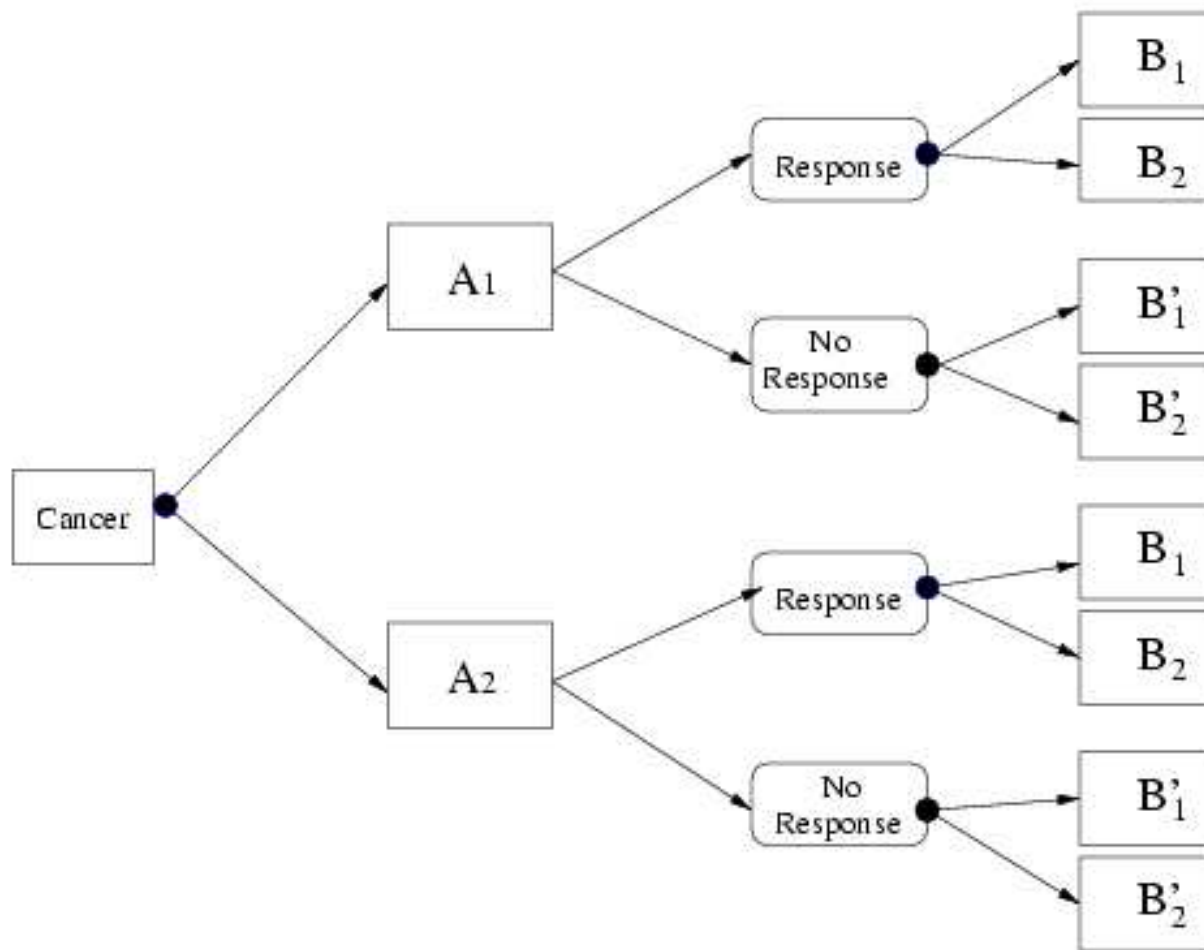
2. Randomized trials for dynamic regimes

Possible ways to compare:

- An *eight-arm* randomized trial?
- Combine information from a *series* of trials?
- *Something else*?

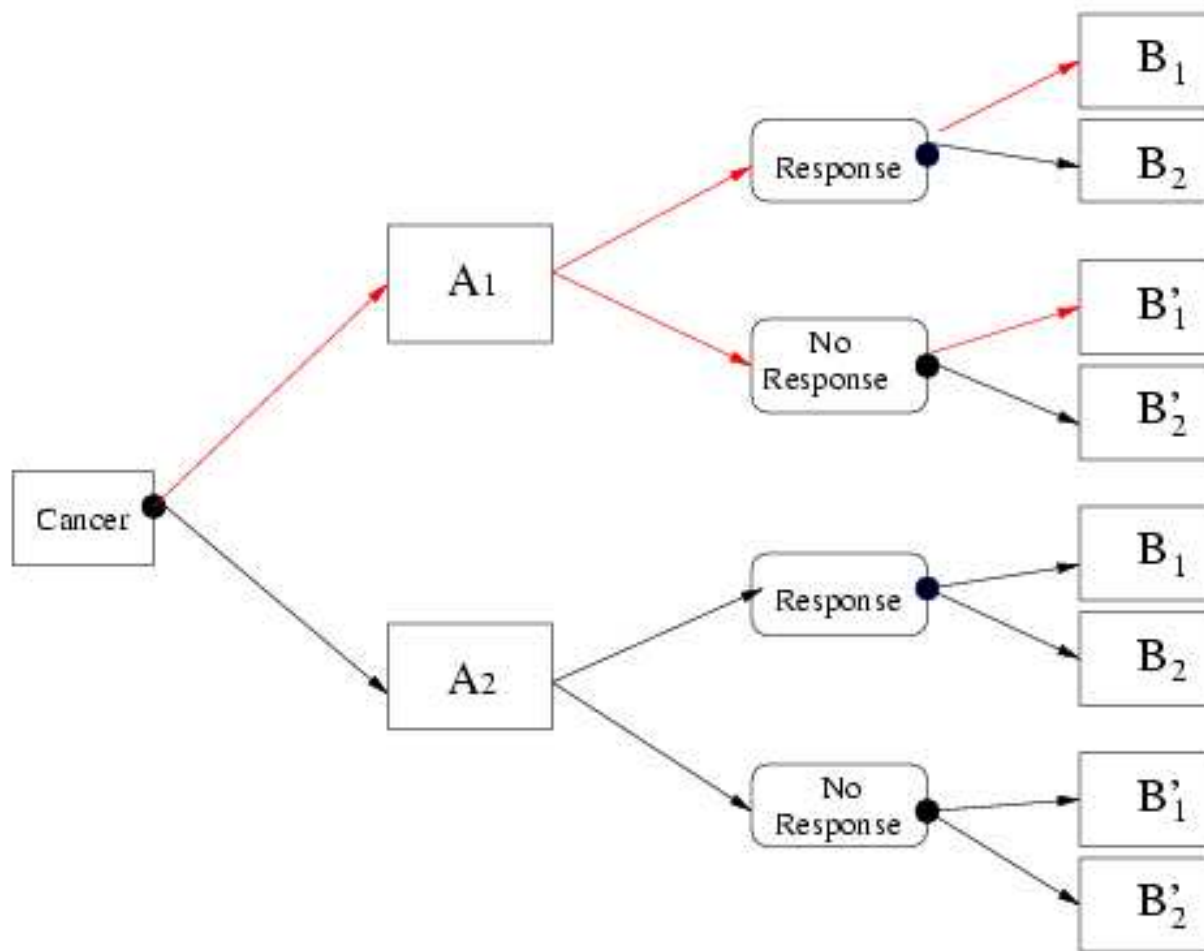
2. Randomized trials for dynamic regimes

“SMART” Trial: Sequential Multiple Assignment Randomized Trial
(Randomization at ●s)



2. Randomized trials for dynamic regimes

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



2. Randomized trials for dynamic regimes

SMART Trials: *Susan* will lay out a rationale and framework for this kind of trial for *designing* and *comparing* dynamic treatment regimes!

- As long as the *number of options* at each “*decision node*” is the same with same probabilities, analysis is *straightforward*
- “*Balanced*”

It turns out: A certain kind of “*not quite as SMART*” trial is common in oncology ...

- Analysis is a little more fancy ...

3. Case study: CALGB 8923

Cancer and Leukemia Group B (CALGB) Protocol 8923: A trial with *two* randomizations, conducted in early 1990s

Background: *Acute myelogenous leukemia (AML)*

- At the time, *standard induction chemotherapy* (daunorubicin+cytarabine)
- Standard chemotherapy \Rightarrow *myelosuppression* \Rightarrow increased risk of death due to *infection* or *bleeding*
- *Add* to standard chemotherapy + *granulocyte-macrophage colony-stimulating factor (GM-CSF)* to reduce risk of these complications (but could possibly *worsen* leukemia...)
- Standard chemotherapy might be followed by “*intensification treatment*” if there is a *response*

3. Case study: CALGB 8923

As before:

- *Step 1* options: A_1 = Standard chemotherapy, A_2 = Standard chemotherapy + GM-CSF
- *If response*, *Step 2* options: B_1, B_2 = “*intensification*” treatments I and II

3. Case study: CALGB 8923

Common oncology trial design: “Two stage randomization”

- After enrollment, *randomize* all subjects to *induction* therapies, e.g., A_1 or A_2 (“*stage 1 randomization*”)
- Observe *intermediate outcome*, e.g., “*response*”
- *Randomize* responding subjects to *maintenance* therapies, e.g., B_1 or B_2 (“*stage 2 randomization*”)
- Subjects not responding *follow up* with their physicians (no “*stage 2*” randomization; *only option*)
- Continue to monitor *all subjects* for the *outcome of interest*, *survival time*
- *Sometimes*: The *nonresponders* are randomized at *stage 2*, *responders* are *not*

3. Case study: CALGB 8923

Four possible regimes:

1. A_1 followed by B_1 if response else follow up = $A_1 B_1$
2. A_1 followed by B_2 if response else follow up = $A_1 B_2$
3. A_2 followed by B_1 if response else follow up = $A_2 B_1$
4. A_2 followed by B_2 if response else follow up = $A_2 B_2$

Question: How do these four regimes *compare* on the basis of *disease-free survival time*?

- E.g., mean disease-free survival time, proportion surviving without disease after 1 year, etc.
- Which regime to recommend?

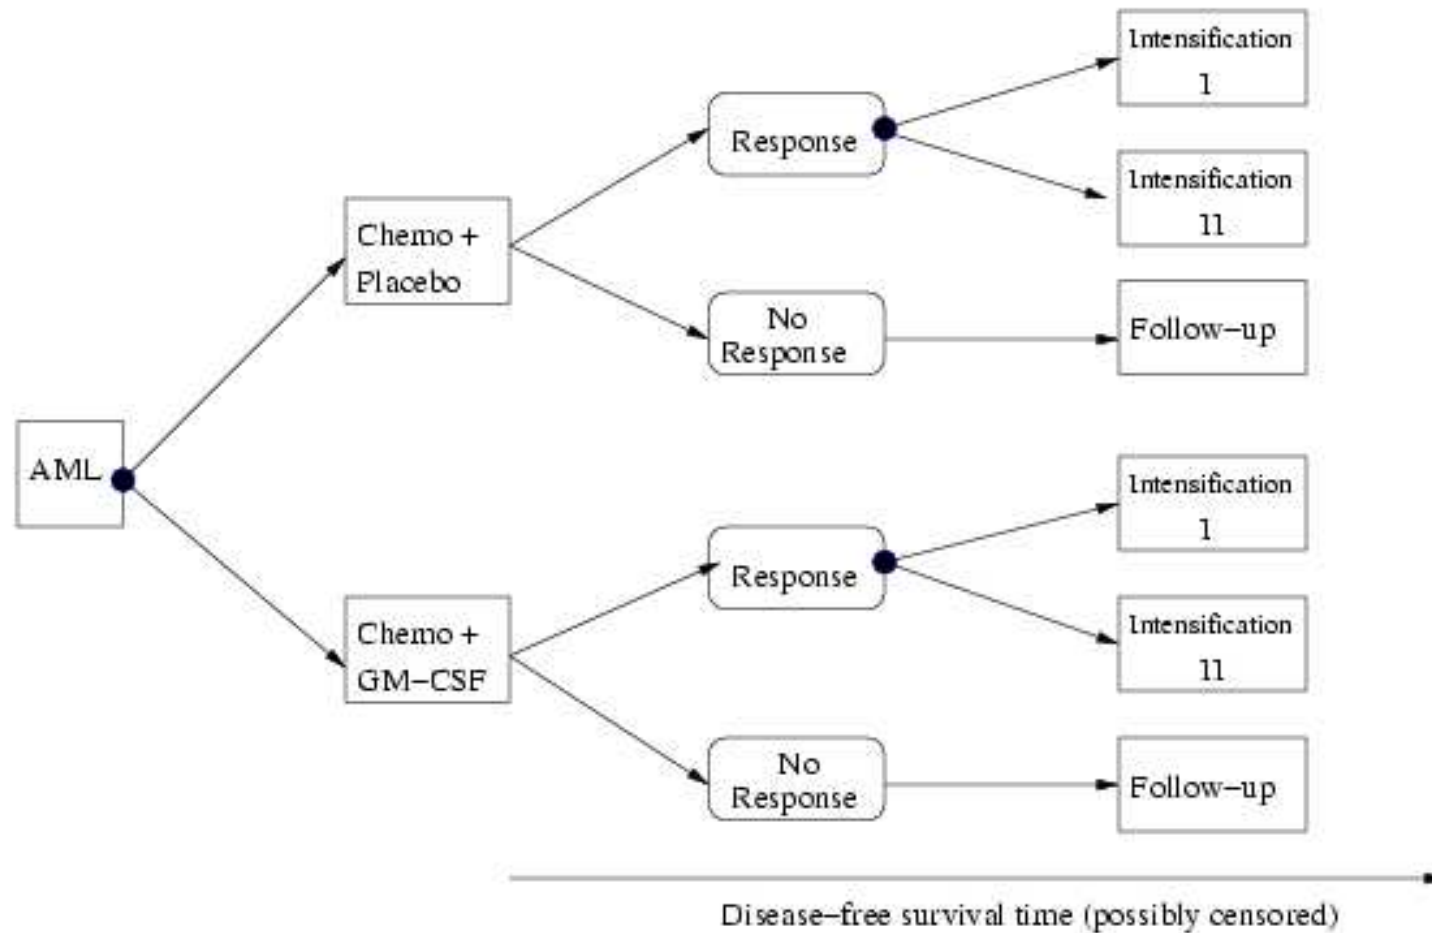
3. Case study: CALGB 8923

CALGB 8923:

- Double-blind, placebo-controlled, two stage randomization trial
- A_1 = standard chemotherapy + *placebo* A_2 = standard chemotherapy + *GM-CSF*
- 338 elderly (> 60 years old) patients with AML
- “*Response*” = complete remission
- B_1, B_2 = intensification treatments I and II
- Goal: Compare the four regimes on the basis of *disease-free survival*

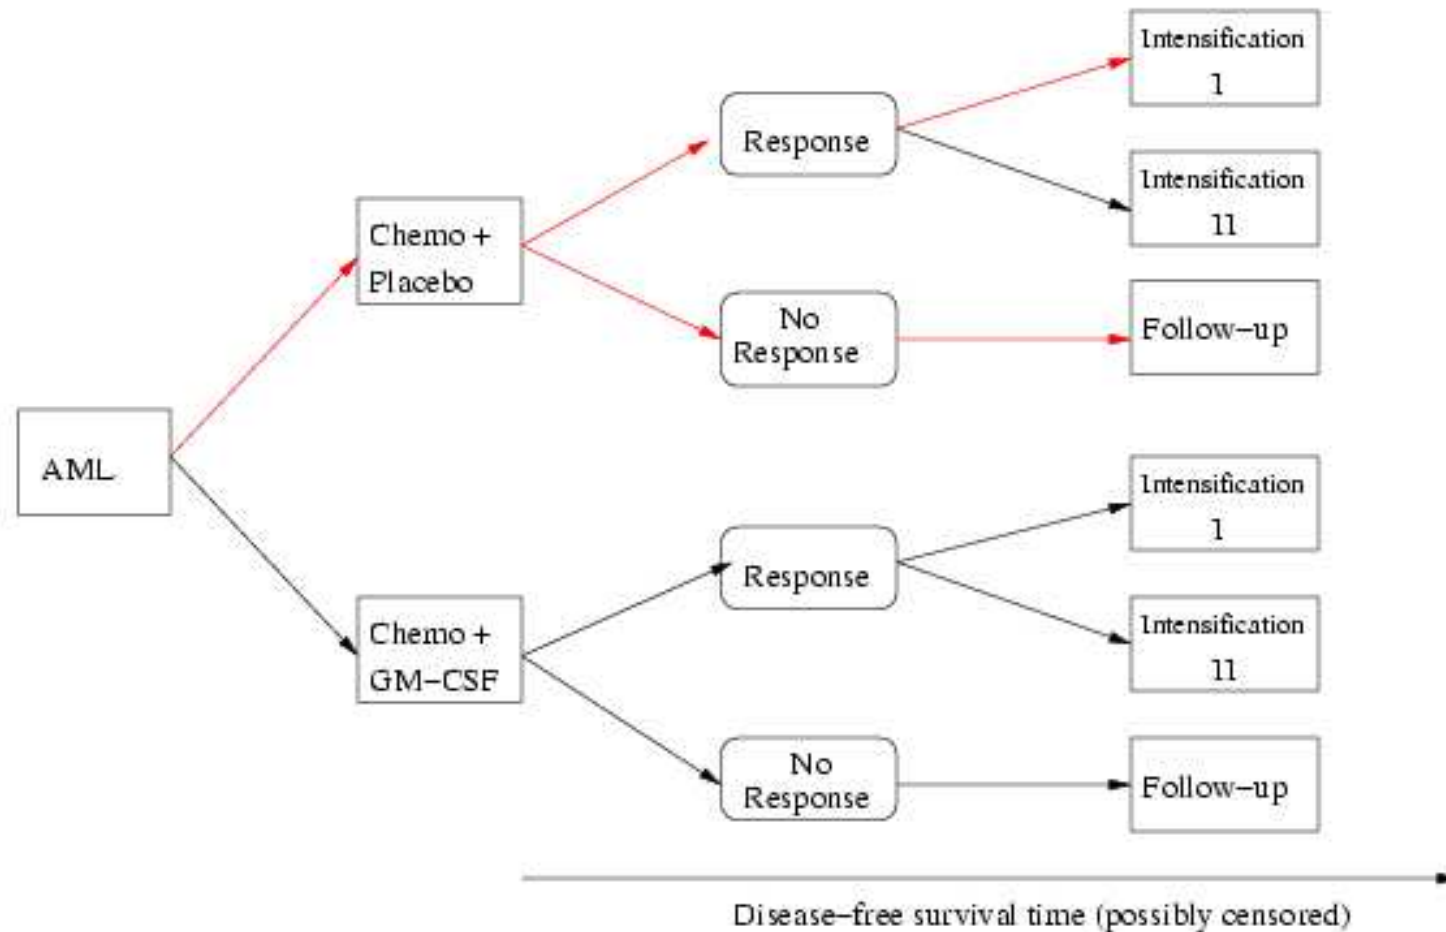
3. Case study: CALGB 8923

Schematic of CALGB 8923: Randomization at ●s



2. Randomized trials for dynamic regimes

Regime A_1B_1 :



4. Analysis

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

Issues:

- Does not address *directly* the question of interest
- An induction therapy (A) may yield *higher proportion of responders* but also have other effects that render subsequent intensification treatments (B) *less effective*
- “*Delayed effects*” (*Susan*)

4. Analysis

Question of interest: For each regime $A_j B_k$, $j = 1, 2$, $k = 1, 2$

- *Estimate* the *mean disease-free survival time* under regime $A_j B_k$
- I.e., estimate mean disease-free survival if the *entire AML population were to follow* regime $A_j B_k$
- “*Following*” $A_j B_k$ means give A_j initially followed by B_k if response else follow up

How to estimate this quantity from the data in the trial?

4. Analysis

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

- *Assume* that whether response occurs depends only on A
- All subjects receiving A_j who *respond* and then are randomized to B_k are *consistent with* $A_j B_k$
- All subjects receiving A_j who *do not respond* and hence are not randomized at stage 2 are *also consistent with* $A_j B_k$
- *Key*: Must *combine* survival times from these subjects in an *appropriate way*...

An appropriate way: This is an “*unbalanced*” SMART trial

- \Rightarrow A *weighted average* of survival times
- Consider this *heuristically*...

4. Analysis

Consider A_1 only (A_2 analogous) : *Ideally*, suppose *everyone* were randomized to A_1B_1

- *Nonresponders* to $A_1 \Rightarrow$ follow up
- *Responders* \Rightarrow all get B_1
- *Natural estimator: Sample average* of *all* survival times (*unweighted*)

In the trial: Suppose *responders* are randomized to B_1 or B_2 with probability $1/2$

- *Nonresponders* to $A_1 \Rightarrow$ follow up (same as before)
- *Half of responders* get B_1 , *half* get B_2
- The half who get B_2 have *missing* survival times as far as A_1B_1 is concerned

4. Analysis

Result: To estimate mean for A_1B_1 from the trial

- The *nonresponders* represent themselves either way \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

4. Analysis

In symbols: Let

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 $I(T_i \leq t)$

5. Wrinkles

Survival outcome: Subjects may *die* before having a chance to respond

- *Nonresponders* at the time of death, $R_i = 0$

Censoring: Survival time may be right-censored at time C_i

- *Assume* $K_1(t) = P(C_i > t | A_1)$
- Consider *restricted survival time*, i.e., survival up to time L such that $K_1(L) > 0$
- *Observe* $V_i = \min(T_i, C_i)$ and $\Delta_i = I(T_i < C_i)$
- If T_i is *not censored* for subject i , $V_i = T_i$, i represents $K_1^{-1}(V_i)$ individuals, including him/herself, who *could have* been uncensored
- *Estimator* becomes

$$n^{-1} \sum_{i=1}^n \frac{\Delta_i Q_i}{K_1(V_i)} V_i$$

5. Wrinkles

Consent of responders: In CALGB 8923, some subjects who *did respond refused to be randomized* at the second stage

- In CALGB 8923, $\sim 90\%$ consent rate among responders
- “*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* else follow up” (so make comparisons without regard to differential rates of consent)
- So *redefine* $\Rightarrow R_i = 1$ if subject i *both* responds *and consents* to further participation
- ... As opposed to attempt to ask the original *causal* question, with this *noncompliance* as a nuisance (\Rightarrow *observational study*)

6. Discussion

Remarks:

- Could equally well *randomize subjects up front* to regimes and use these same estimators
- *Fancier* (in terms of efficiency) estimators are possible
- Methods for *testing* also possible
- If SMART trial is *balanced*, *no need* to do *weighting*

Looking forward to Susan:

- Dynamic treatment regimes are what is done in *clinical practice*
- The regimes here are simple and preconceived: *two stages* only, *decision rule* at step 2 based on the single variable “*response*”
- Methods to *design* dynamic treatment regimes are needed

7. Demonstration using potential outcomes

One way to formalize the rationale for weighting: Again consider A_1 regimes only (A_2 analogous)

- Suppose there are n subjects randomized to A_1 and that subject i has *potential outcomes* T_{11i}, T_{12i}
- T_{1ki} = survival time i *would have* if i *were to follow* (*or be offered*) $A_1 B_k$, $k = 1, 2$

Question of interest: Estimate mean disease-free survival if the *entire AML population were to follow* regime $A_1 B_k$

- Distributions of the T_{1k} represent survival in the population *if all subjects* followed $A_1 B_k$, $k = 1, 2$
- \Rightarrow Want to estimate $\mu_{1k} = E(T_{1ki})$

7. Demonstration using potential outcomes

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

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

- $R_i = 1$ if i *responds*, $R_i = 0$ if not
- $Z_i = k$ if i is *randomized at stage 2* to B_k , $k = 1, 2$
(defined only if $R_i = 1$)
- $P(Z_i = 1 \mid R_i = 1) = \pi$ = probability of second stage randomization to B_1 (after first stage randomization to A_1) *if response*

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

- The estimators discussed (based on *observed* data) may be shown to be *consistent* for μ_{11} , e.g., $n^{-1} \sum_{i=1}^n Q_i T_i$

7. Demonstration using potential outcomes

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

Assume: For subjects randomized to A_1

- If $R_i = 0$, T_{11i} and T_{12i} are the same; thus

$$T_i = (1 - R_i)T_{11i} + R_i I(Z_i = 1)T_{11i} + R_i I(Z_i = 2)T_{12i}$$

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 want to show

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

7. Demonstration using potential outcomes

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

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 $Q_i = 1 - R_i + R_i I(Z_i = 2)(1 - \pi)^{-1}$,

References

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.

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*, in press.

These slides available at:

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