

Toward Individualizing Treatment to the Patient: 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, and why study them?
2. Clinical trials to study dynamic treatment regimes
3. Studying dynamic treatment regimes based on observational data
4. Constructing dynamic treatment regimes
5. Concluding remarks

The thought leaders:

- Jamie Robins
- Susan Murphy
- Many others!

1. What is a dynamic treatment regime?

Clinical practice: Treatment of *chronic disease* is an *ongoing process*

- Clinicians *do not* make a *once-and-for-all* decision
- Rather, clinicians *manage* a patient's illness
- Clinicians routinely *adjust*, *change*, *add*, or *discontinue* treatment based on *progress*, *side effects*, *patient burden*, *compliance*, etc.
- *Among-* and *within-* patient *heterogeneity*
- “*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*, *patient preference*
- *Practice guidelines* based on *clinical evidence* and *expert opinion*, e.g., NIH guidelines for treatment of HIV infection (<http://www.aidsinfo.nih.gov/>); ACC/AHA Practice Guidelines (<http://www.americanheart.org/presenter.jhtml?identifier=2158>)

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
- Taken together, these *systematic rules* define an *algorithm* for treating a patient that seeks to *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 (variables)* on the patient *up to that point*...
- ...based on a *decision rule* that takes these *variables* as *input* and *outputs* the *next treatment step* for the patient
- An *algorithm* that dictates how treatment of a patient should proceed over time

A simple 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 *maintain or intensify* 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*


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*” = complete or partial remission, extent 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?

Obviously: This is a *very simple* regime


- Involves a *single decision rule* that takes as *input* the *single, binary variable* “*response*”
- Clearly, clinicians take into account *much more information* in making decisions, and they make *multiple decisions* in order to “*individualize*” treatment
- *Ultimate objective*: Define regimes with more *complex decision rules* that translate *many variables* into treatment decisions...
- ...that mimic more closely *clinical practice*

1. What is a dynamic treatment regime?

Themes for this talk:

- Although clinicians engage in *sequential treatment decision-making*, we do not *study* this!
- Can we formally *study* and *compare* different regimes on the basis of *outcome*?
- In fact, can we *design clinical trials* to do so?
- Can we identify *statistical methods* for analysis of the data?
- In fact, can we use *data* from such trials and *observational studies* to identify “*good*,” more *realistic* regimes that get closer to *individualizing* treatment?

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' ”



(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*

1. What is a dynamic treatment regime?

Important, part 3: *Do not confuse* dynamic treatment regimes *themselves* with *response-adaptive designs* for studies of traditional treatments

- A dynamic treatment regime is an *algorithm* for treating *a single patient* that takes as *input data on that patient only*
- This has *nothing to do* with *other patients* in a study

2. Clinical trials for studying treatment regimes

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

- Consider our *cancer example*
- 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. Clinical trials for studying 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_1 if response, else B'_2
7. A_2 followed by B_2 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. Clinical trials for studying 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. Clinical trials for studying treatment regimes

One 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* data, where *treatments actually received* over time have been recorded (with other information) for each subject

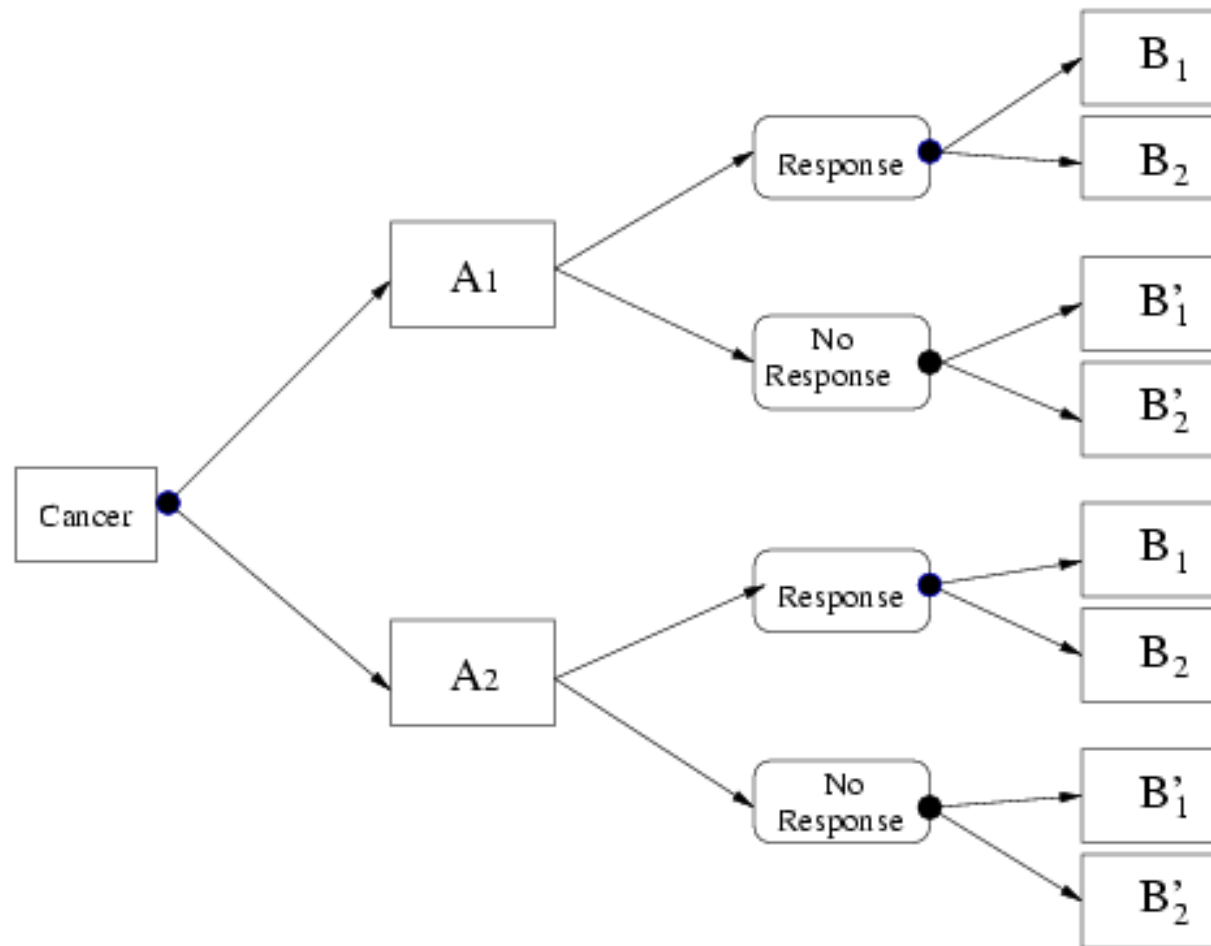
2. Clinical trials for studying treatment regimes

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?

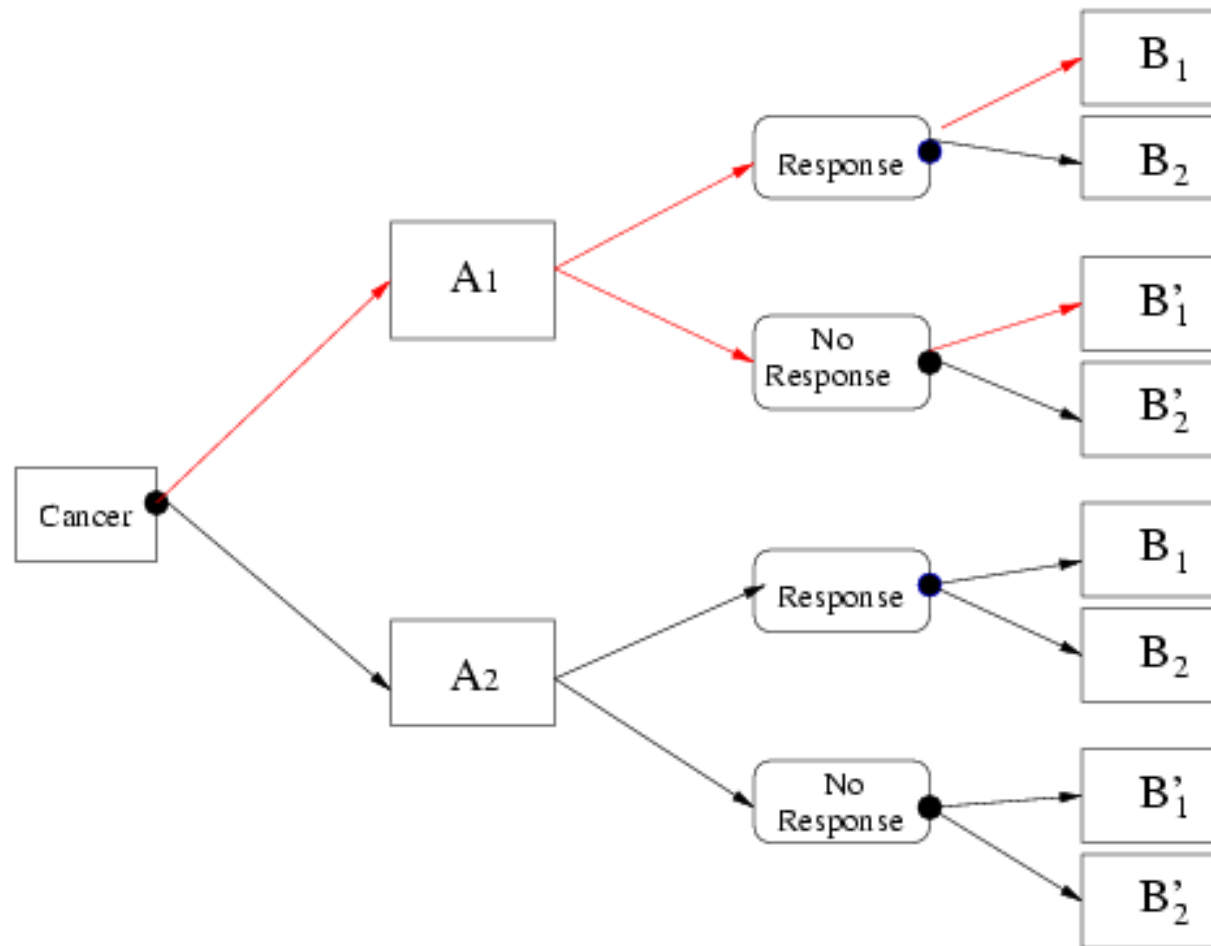
2. Clinical trials for studying treatment regimes

“SMART:” Sequential Multiple Assignment Randomized Trial, e.g., Lavori and Dawson (2004) (Randomization at ●s)



2. Clinical trials for studying treatment regimes

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



2. Clinical trials for studying treatment regimes

Examples of SMARTs:

- CATIE (2001) Treatment of psychosis in Alzheimer's, schizophrenia
- STAR*D (2003) Treatment of Depression

Remarks:

- There is really no *conceptual difference* between randomizing *up-front* or *sequentially*
- Each has *advantages* and *disadvantages*, e.g., *consent*, *balance*
- What is *important*: Making best use of the data for *analysis*...

Reference:

Murphy, SA. (2005). An experimental design for the development of adaptive treatment strategies, *Statistics in Medicine* **24**, 1455–1481.

2. Clinical trials for studying 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. Clinical trials for studying 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...

2. Clinical trials for studying treatment regimes

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*

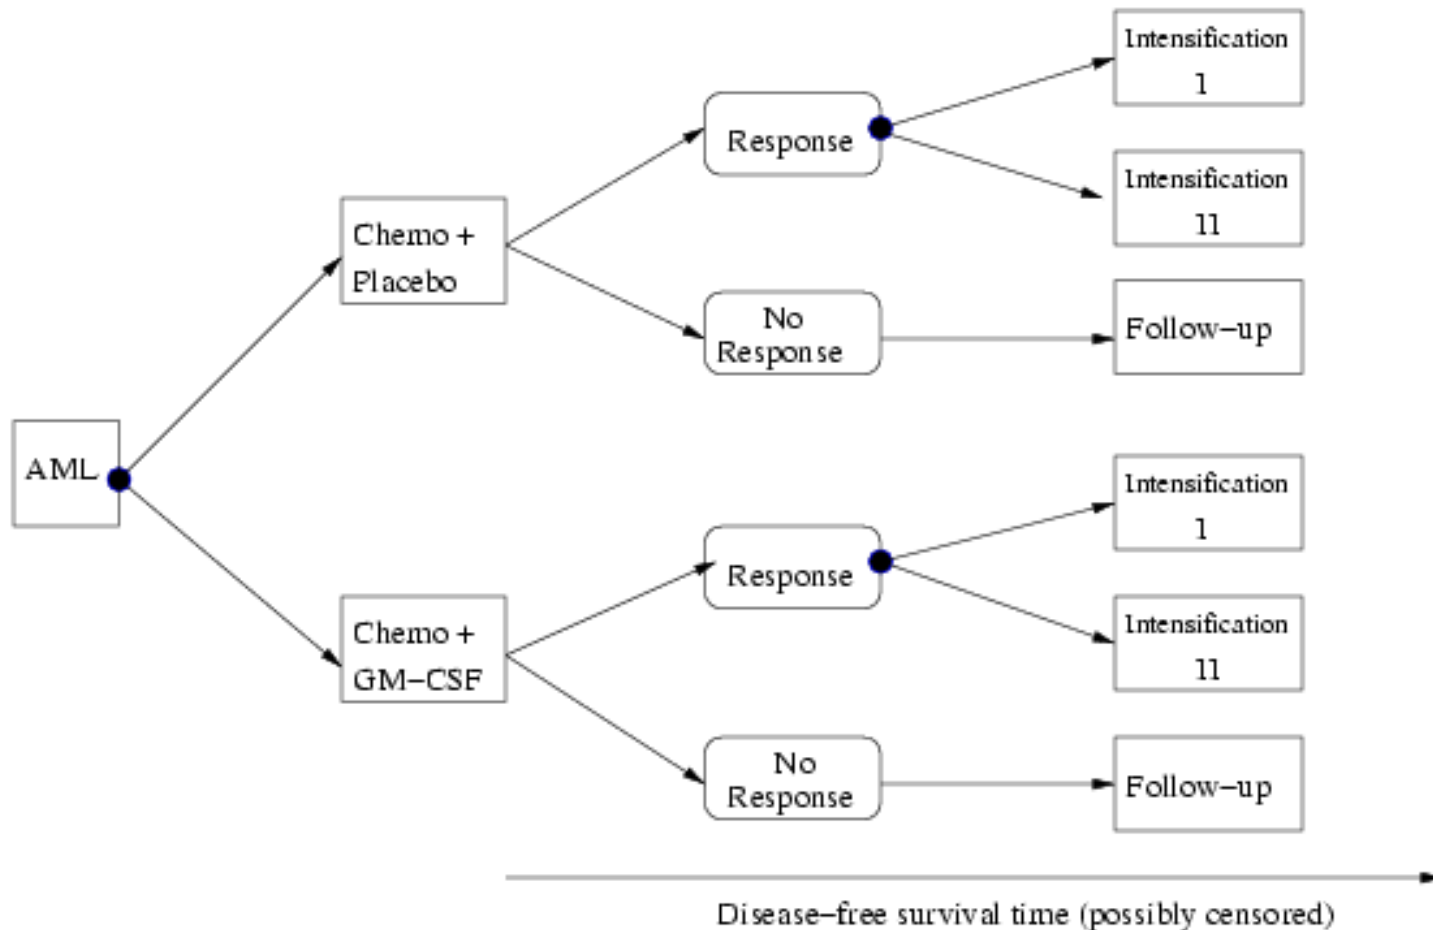
2. Clinical trials for studying treatment regimes

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

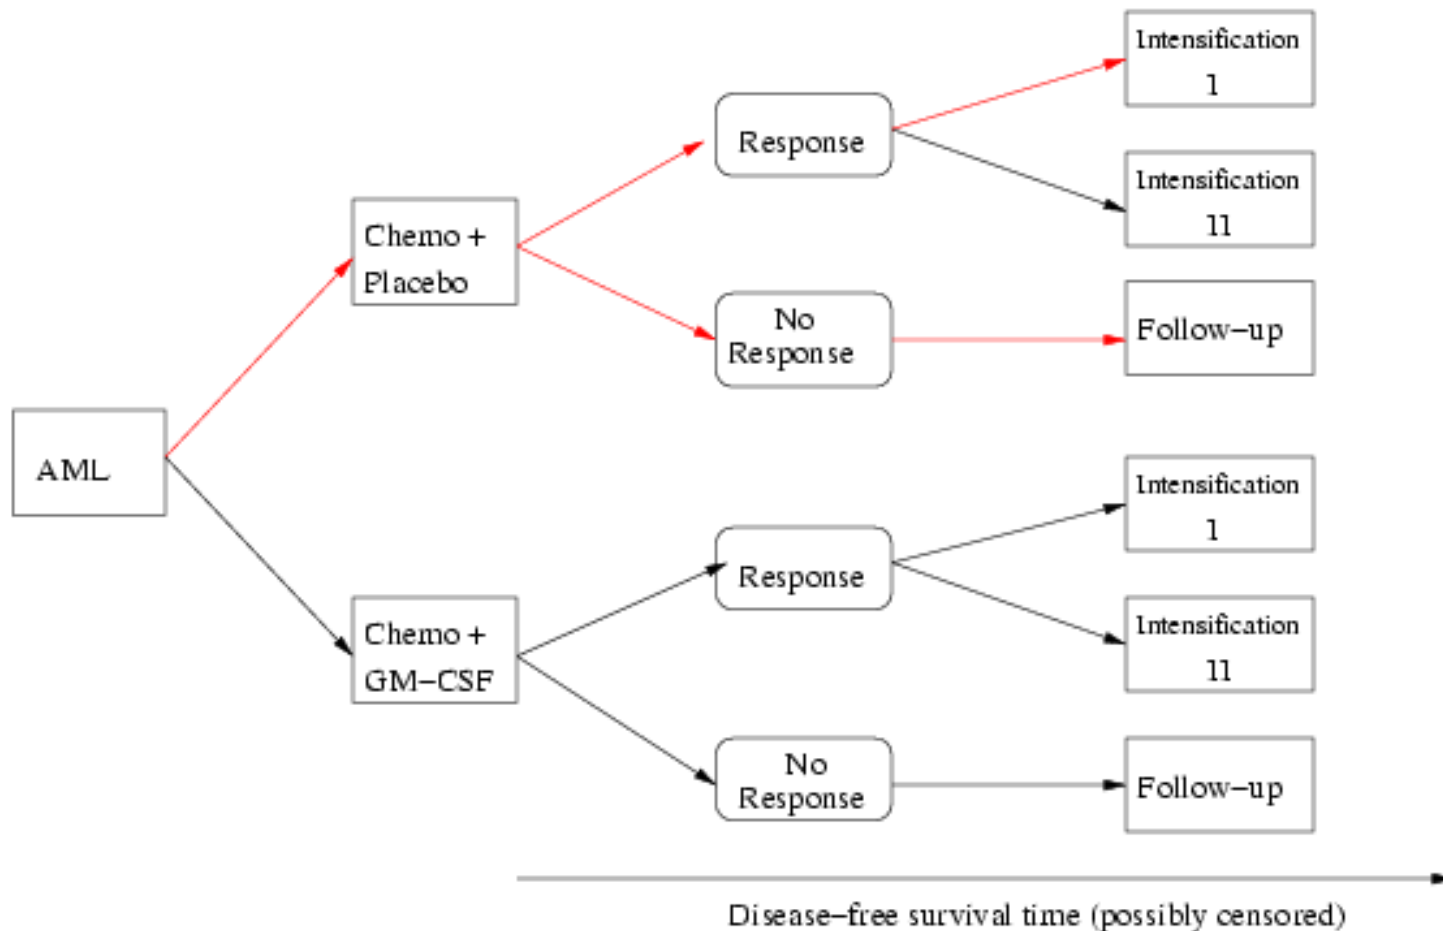
2. Clinical trials for studying treatment regimes

Schematic of CALGB 8923: Randomization at ●s



2. Clinical trials for studying treatment regimes

Regime A_1B_1 :



2. Clinical trials for studying treatment regimes

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*

2. Clinical trials for studying treatment regimes

Analysis to compare dynamic treatment regimes:

- Estimate *mean survival* if the *entire AML population* were to follow each regime $A_j B_k$, $j = 1, 2$, $k = 1, 2$
- *Illustrate* with regime $A_1 B_1$

Basic idea: To *estimate the mean* for $A_1 B_1$, use data from all subjects whose *realized experience* is *consistent with* having followed $A_1 B_1$

- $A_1 \Rightarrow$ *response* $\Rightarrow B_1$
- $A_1 \Rightarrow$ *no response* \Rightarrow follow up with physician
- *Combine* survival times from these subjects in an *appropriate way*...

2. Clinical trials for studying treatment regimes

Consider A_1 subjects: 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
- To estimate *mean survival* for A_1B_1 , use a *weighted average*
- The *nonresponders* represent themselves \Rightarrow weight = 1
- Each *responder* who got B_1 represents him/herself and *another similar subject* who got randomized to $B_2 \Rightarrow$ weight = 2
- 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

2. Clinical trials for studying treatment regimes

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 $I(T_i \leq t)$

2. Clinical trials for studying treatment regimes

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_1 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*)

2. Clinical trials for studying treatment regimes

Remarks:

- The *key* to the validity of the *weighted analysis* is the *randomization* at each step \Rightarrow subjects are prognostically similar
- This idea can be *extended* to any number of steps and numbers of options at each step

2. Clinical trials for studying treatment regimes

Issues in design of SMARTs:

- Identify *decision points* subsequent to the initial step of treatment where *key treatment decisions* would be made
- At these decision points, there should be *more than one treatment option* and a *lack of agreement* and *uncertainty* regarding the choice of options
- These considerations will define a set of *regimes*
- *Keep it simple* – small number of *feasible options* and small number of *steps*
- *Keep it simple* – *decision rules* taking as input *simple*, *low-dimensional* variables

2. Clinical trials for studying treatment regimes

Why do this?

- Even if primary interest is in A_1 vs. A_2 , we can still do the usual *intention-to-treat* analysis
- But, by *systematizing* part of what comes next, we are in a position to *learn more* from the data
- In fact, a *major reason* to carry out SMARTs is to *inform development of regimes* that get closer to *individualizing treatment*
- More shortly...

3. Studying regimes based on observational data

Often: Sources of data are *observational* in nature

- Databases from *prospective studies*, *registries*
- Databases from *completed clinical trials* – questions on issues *other than* the *randomized treatments*
- What are “*good*” strategies for using *antiretroviral* (ARV) therapies to treat HIV-infected subjects *over time*?
- Should one *switch treatments*, and on *what basis*?
- If treatment is effective, *what should be the duration*?
- Should *transfusion* be given, and *under what conditions*?

Questions about dynamic treatment regimes...

3. Studying regimes based on observational data

For example: ESPRIT trial – Integrilin vs. placebo in PCI/stent patients

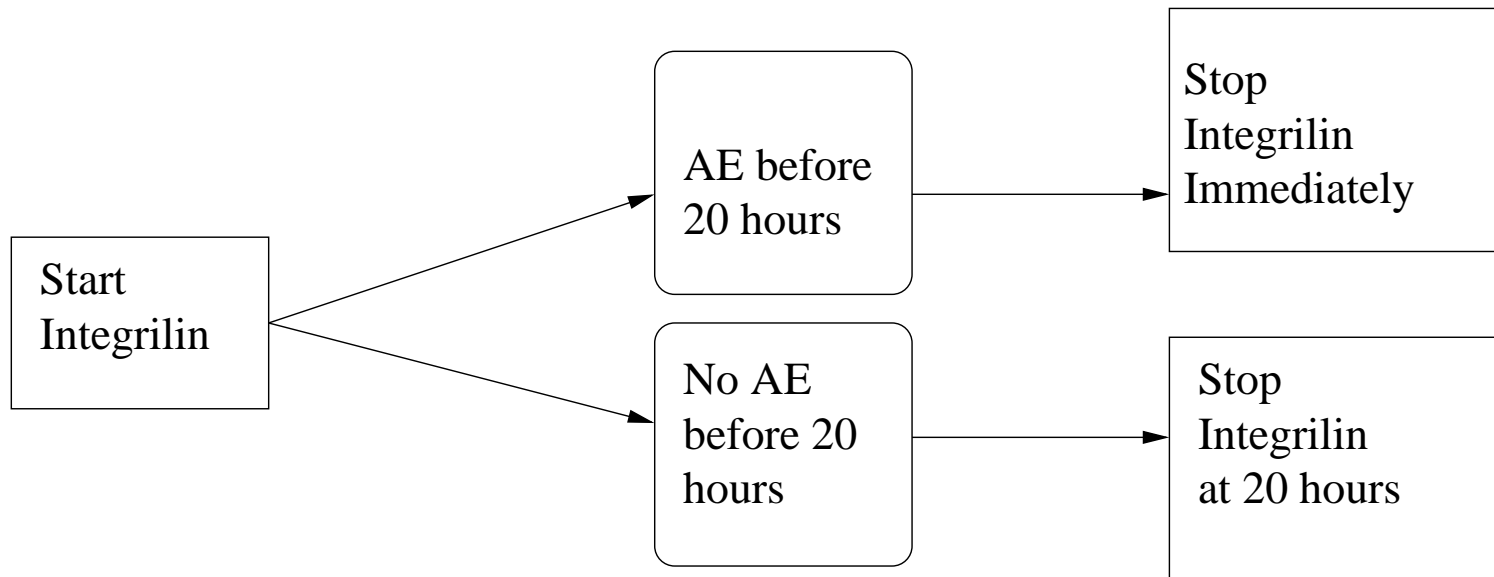
- Duration of Integrilin infusion was left to *physician discretion*
- *Protocol*: Duration of 18 – 24 hours with *mandatory discontinuation for adverse events*
- The *primary analysis* showed Integrilin *superior to* placebo on the basis of a composite outcome
- \Rightarrow Attention focused on determining the “*recommended*” *treatment duration*

More precisely: *Treatment duration of t hours* means *infuse for t hours or until a treatment-terminating adverse event, whichever comes first*

- This is an example of a *dynamic treatment regime*, because realized duration depends on the *variable* adverse event status

3. Studying regimes based on observational data

Duration regime of 20 hours:



3. Studying regimes based on observational data

Issue: Data are *observational* with respect to the questions

- *Decisions* on *duration of treatment*, on *switching treatments*, on *stopping/starting ARV therapy*, etc, were *not randomized*
- Made at *physician/patient discretion*

Difficulties for studying regimes in this setting:

- *Confounding* – subjects receiving one treatment or another *may not be prognostically similar*
- E.g., subjects who *switched treatments* may be *sicker*, *older*, etc
- Standard methods are available to *adjust for confounding*, e.g., *regression*, *propensity scores*, etc, ...
- ... *however*, the *time-dependent nature* of treatment causes *additional complications* that standard methods *cannot handle*

3. Studying regimes based on observational data

The problem: In an *observational study*

- Receipt of treatment *over time* may depend on *intermediate variables*
- E.g., a clinician may place an HIV patient *back on ARV therapy* because his *viral load* has risen
- *Temptation*: “Adjust” for such *time-dependent confounding*
- E.g., one might adopt a Cox regression model for outcome including *time-dependent intermediate variables and treatments*
- *However*: Part of the effect of treatment on outcome may be *mediated through* intermediate variables
- \Rightarrow Adjustment would *incorrectly remove* this effect and hence *misrepresent* the true treatment effect

3. Studying regimes based on observational data

How to resolve this?

- Relate this setting back to the SMART setting. . .
- Requires a *critical* but *unverifiable assumption*

Sequential randomization assumption: *At any point where a treatment decision is made, the treatment received (among the options available) depends only on the past history of the patient and not additionally on his/her future prognosis*

- *History*: past treatments, baseline and intermediate variables
- At some level, this *must be true*
- In a SMART, this is *automatic through randomization*
- *However*, in an *observational study* this assumption is tenable *only if all information used to make decisions is collected*

3. Studying regimes based on observational data

Under the sequential randomization assumption: To make inference on *dynamic treatment regimes*

- Can use *weighted methods* similar to those discussed earlier
- *Critical difference*: Rather than weighting based on *known randomization probabilities*, weighting is based on the *propensities of receiving treatment* as a function of past history
- Requires *modeling and estimation* of the propensities

Jamie Robins: Has pioneered

- The *statistical framework* that formalizes all of this
- Associated *methods* for inference on dynamic treatment regimes

4. Constructing dynamic treatment regimes

Can we construct “better” regimes? From SMARTs

- SMARTs will generally study and compare *simple regimes*
- But the *ultimate goal* is to develop regimes that get closer to *individualizing treatment*
- If *intermediate information* is collected at each step, one is in a position to identify “*tailoring variables*” at each step
- That is, variables that *affect outcome differentially* by treatment at that step. . .
- . . . and hence should be incorporated in the *decision rule* at that step
- *Randomization* at each step is the key

Susan Murphy: Has pioneered *using SMARTs* this way

4. Constructing dynamic treatment regimes

Methodological challenges:

- Methods for doing this must incorporate the *effect of future treatment decisions* when evaluating the *present treatment decision*
- Such methods have been developed by *statisticians*, *computer scientists*, and others and need to be adapted to this setting
- *Susan Murphy*, *Jamie Robins*, and colleagues have pioneered *statistical methods* for inferring the *optimal regime* from either type of data
- *Computer scientists* have developed parallel *reinforcement learning* techniques

4. Constructing dynamic treatment regimes

Another approach: Exploit *mechanistic models*

Structured Treatment Interruption (STI) for acute HIV-1 infection:

- Potent ARV therapy *cannot be taken continually*
- Side effects, burden, cost, drug resistance, ...
- \Rightarrow Cycles of therapy followed by *interruption*
- When to *interrupt*? When to *re-initiate*? On what *basis*?

HIV dynamic models: Formalize hypotheses about *interplay* between *HIV* and *immune system* happening *within a subject*

4. Constructing dynamic treatment regimes

Nonlinear dynamical system: Ordinary differential equations, $\dot{U} = \frac{dU}{dt}$

$$\dot{T}_1 = \lambda_1 - d_1 T_1 - \{1 - \epsilon_1 u(t)\} k_1 V_I T_1$$

$$\dot{T}_2 = \lambda_2 - d_2 T_2 - \{1 - f \epsilon_1 u(t)\} k_2 V_I T_2$$

$$\dot{T}_1^* = \{1 - \epsilon_1 u(t)\} k_1 V_I T_1 - \delta T_1^* - m_2 E T_1^*$$

$$\dot{T}_2^* = \{1 - f \epsilon_1 u(t)\} k_2 V_I T_2 - \delta T_2^* - m_2 E T_2^*$$

$$\begin{aligned} \dot{V}_I = & \{1 - \epsilon_2 u(t)\} 10^3 N_T \delta (T_1^* + T_2^*) - c V_I - \{1 - \epsilon_1 u(t)\} \rho_1 10^3 k_1 T_1 V_I \\ & - \{1 - f \epsilon_1 u(t)\} \rho_2 10^3 k_2 T_2 V_I \end{aligned}$$

$$\dot{V}_{NI} = \epsilon_2 u(t) 10^3 N_T \delta (T_1^* + T_2^*) - c V_{NI}$$

$$\dot{E} = \lambda_E + \frac{b_E (T_1^* + T_2^*)}{(T_1^* + T_2^*) + K_b} E - \frac{d_E (T_1^* + T_2^*)}{(T_1^* + T_2^*) + K_d} E - \delta_E E$$


+ *initial conditions* $\{T_1(0), T_2(0), T_1^*(0), T_2^*(0), V_I(0), V_{NI}(0)\}$

Treatment input: $u(t) = 1$ if therapy given at time t , $= 0$ if not

$$\Rightarrow \text{CD4 count} = T_1 + T_1^*, \quad \text{viral load} = V_I + V_{NI}$$

4. Constructing dynamic treatment regimes

Main measures: *CD4 T-cell count* and *viral load*



4. Constructing dynamic treatment regimes

Statistical model:

- Each subject has his/her own *dynamic parameters* $\lambda_1, k_1, \epsilon_1, \lambda_2, k_2, \epsilon_2, c, \delta, m_1, m_2$ etc.
- These *vary* in the population according to some distributions
- *Observations* on only some of the model states, subject to *variation*
- *Hierarchical nonlinear model*

Challenge: Can this *mathematical/statistical* model be used to develop *dynamic regimes* to recommend to a *population*?

- I.e., determine the *input function* $u(t)$
- *Control theory*: Mathematical theory for *modifying* the behavior of dynamical systems through *control of system inputs*
- Use *feedback control* methods to design *dynamic STI regimes*

4. Constructing dynamic treatment regimes

The test: Collaboration with *H.T. Banks* (applied mathematician at NCSU) and *Eric Rosenberg* (immunologist/infectious disease physician at Mass General Hospital) to develop methods to do this

- More *realistic HIV dynamic models* informed by intensive *longitudinal data* from Eric's *acute infection cohort*
- Devising *control theoretic methods* that yield *practically feasible* rules and inputs
- *Statistical-model-based simulation* to determine effect on population
- Run a *clinical trial* to evaluate treatment strategies suggested by the models

5. Concluding remarks

- *Dynamic treatment regimes* are *algorithms* that *operationalize* how clinicians *practice* medicine
- *Methods for comparing regimes* based on *clinical trials* (e.g., SMARTs) and *observational studies* are available
- If all baseline and intermediate information that may be important for *treatment decision-making* has been collected, can make progress on developing *more refined regimes* that get closer to *individualizing treatment* (still *methodological work* to be done!)
- *Statisticians* and *clinicians* should be open to thinking about treatment as *sequential multi-stage decision-making* and studying *dynamic treatment regimes*

Some references

- Johnson BA, Tsiatis AA. (2004). Estimating mean response as a function of treatment duration in an observational study, where duration may be informatively censored. *Biometrics* **60**, 315–323.
- Johnson BA, Tsiatis AA. (2005). Semiparametric inference in observational duration-response studies, with duration possibly right-censored. *Biometrika* **92**, 605–618.
- Lavori P, Dawson R. (2004). Dynamic treatment regimes: practical design considerations. *Clinical Trials* **1**, 9–20.
- 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. (2003). Optimal dynamic treatment regimes (with discussion). *JRSS-B* **65**, 331-366.
- 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>
- Murphy SA, van der Laan MJ, Robins JM. (2001). Marginal mean models for dynamic regimes. *JASA* **96**, 1410-1423.

Some references

- 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
- 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.
- Robins JM. – Visit <http://www.hsph.harvard.edu/robins/> Even
- 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.

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