Machine Learning Techniques to Personalize the Medicine

PhUSE Delhi 2018
Vinodhkumar Katikala, GSK

Agenda

- What is Personalized Medicine?
- Dynamic Treatment Regimens
- The Big Statistical Questions
- Clinical Reinforcement trials:
  - ADHD SMART Example
  - Q-Learning
  - PROC QLEARN

What is Personalized Medicine?

In general: One size does not fit all

- Multiple treatment options may be available
- Patient heterogeneity
  - Across patients: What works for one patient may not work for another
  - Within patients: What works now may not work later

Premise: Use information on a patient's characteristics to determine which treatment option she/he should receive (and when. . .)

- Genetic/genomic, demographic,.
- Physiologic/clinical measures, medical history,. . .

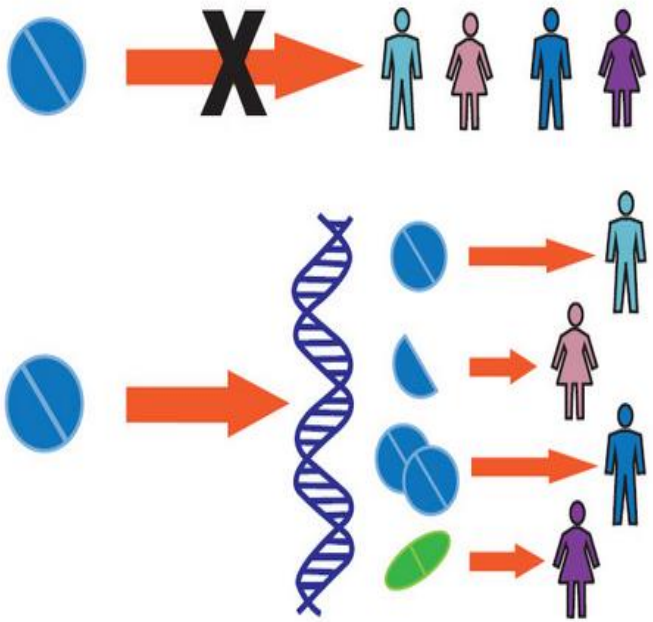

Dynamic Treatment Regimens

A **Dynamic Treatment Regimen (DTR)*** is a **sequence of decision rules**, one per stage of clinical intervention, that specify how to adapt the **type, dosage and timing of treatment** according to individual patient's ongoing **response, adherence, burden, side effects, preference, and past treatments.**

- Each decision rule takes a patient's treatment and covariate history as inputs, and outputs a recommended treatment

Simple example:

Which treatment to give patients who present with **primary operable breast cancer ?**

Options: L-phenylalnanine mustard and 5-flourouracil (PF) or PF + Tamoxifen (PFT)

If age < 50 yrs and progesterone receptor < 10 fmol ⇒ PF (1); else ⇒ PFT (0)

* *Alternatively known as* ***adaptive treatment strategies*** *or* ***adaptive interventions*** *or* ***treatment policies*** *or* ***stepped care models*** *or* ***expert systems***

The Big Statistical Questions

1. What is the right kind of data for comparing two or more DTRs, or estimating optimal DTRs? What is the appropriate study design?
   - **Clinical Reinforcement Trials [S**equential **M**ultiple **A**ssignment **R**andomized **T**rial (**SMART**)]

2. How can we compare pre-conceived, embedded DTRs?
   - Primary analysis of SMART data

3. How can we estimate the “optimal” DTR for a given patient and how to make it **evidence-based** or **data-driven**?
   - secondary analysis of SMART data
   - Machine Learning techniques [e.g. Q-learning, a stage wise regression-based approach]

Clinical Reinforcement trials:

The basic process of reinforcement learning involves

– Trying a sequence of actions and recording the consequences of those actions, statistically estimating the relationship between actions and consequences, and then choosing the action that results in the most desirable consequence.

These can be done by **SMART** concept

– Each patient is randomized at each decision time to a possibly continuous range of treatment possibilities (drug, dose etc.).

– At the end of the first stage, reinforcement learning is used to estimate optimal treatment per prognostic values at each decision time.

– The first stage is followed by a second, confirmatory stage.

Example: SMART trial

– SMART designed to build an adaptive intervention to optimize school performance of children with Attention Deficit Hyperactivity Disorder (ADHD)

The data used here is simulated.

Primary outcome: Teacher-rated Impairment Rating Scale (TIRS)

The data includes:

Baseline covariates or other measures of Stage I (O1): Race, Prior medication, ODD (Oppositional Defiant Disorder) Diagnosis, ADHD Symptoms

Stage I treatment(A1): MED vs BMOD(Behavioral Modification Treatment)

Stage II covariates(O2): Adherence to first stage treatment

Stage II treatment: intensified version of the first-stage treatment or augmentation of the first-stage treatment with the other treatment

Example: ADHD(Attention-deficit hyperactivity disorder)

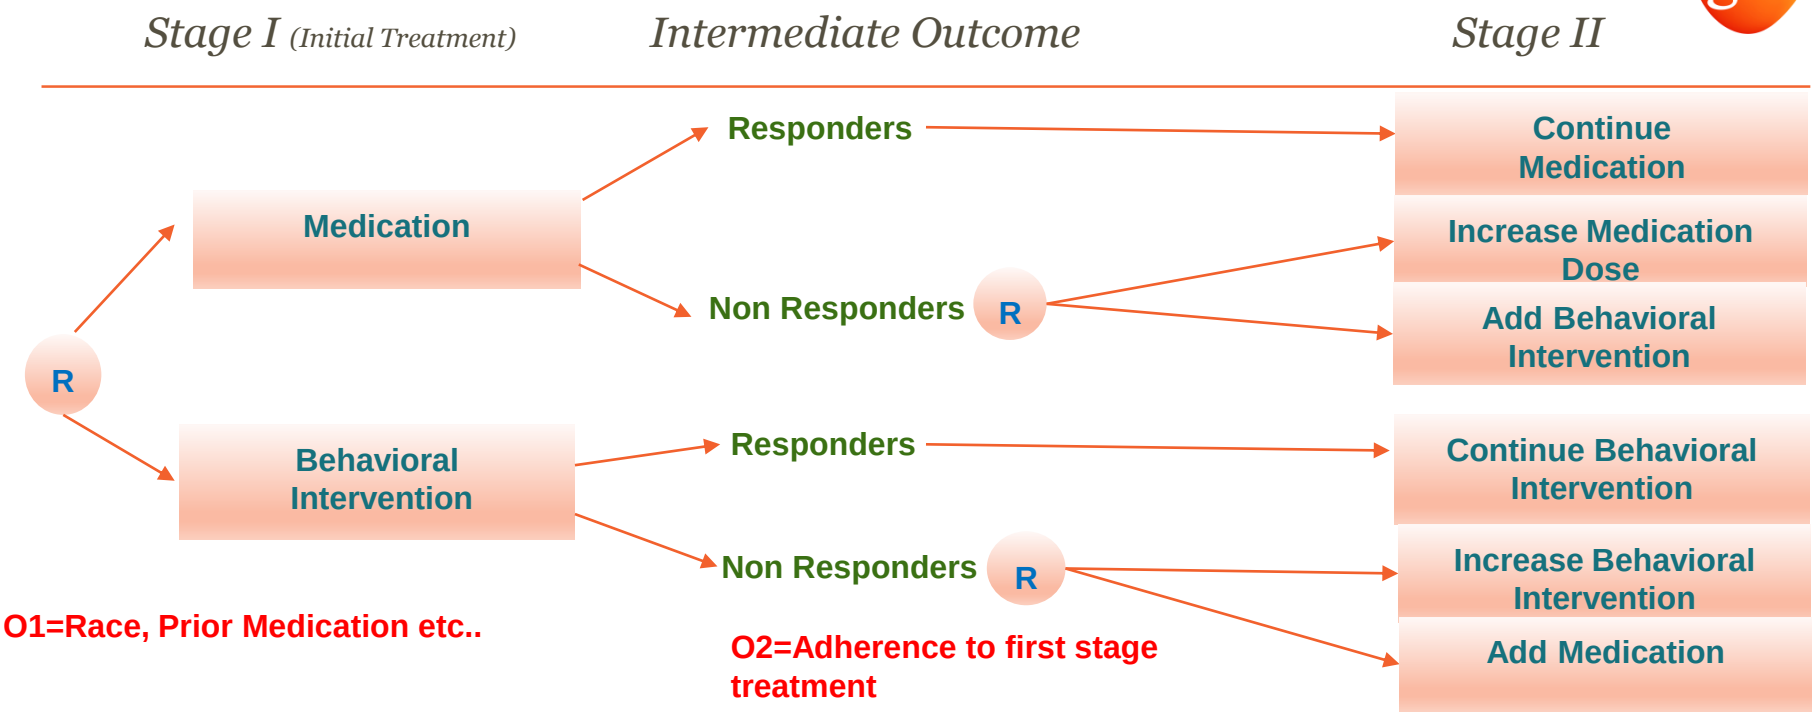

O1 : Other measures or Covariates of stage I
O2 : Other measures or Covariates of stage II

Example: ADHD(Attention-deficit hyperactivity disorder)

ID	o11c	o12c	o13c	o14c	a1	r	o21c	o22c	a2	Y
5	0	-0.70404	1	0	-1	1		0		4
9	1	-1.01254	1	1	-1	1		0		4
23	0	0.14336	1	1	-1	1		1		5
24	1	1.07187	0	1	-1	1		0		1
27	0	-0.84515	0	0	-1	1		1		4
30	0	-0.21032	0	1	1	0	7	0	1	1

O11c : ODD (Oppositional Defiant Disorder) Diagnosis: whether the child was (coded as 1) or was not (coded as 0) diagnosed with ODD before the first-stage treatment.

O12c : ADHD symptoms at the end of the previous school year (larger values reflect fewer symptoms)

O13c : Medication prior to first-stage treatment(1=received, 0=not received)

O14c : Race variable: white (coded 1) versus non-white (coded 0).

a1 : Stage I treatment(1=Medication ,-1=Behavioral Intervention)

Example: ADHD(Attention-deficit hyperactivity disorder)

ID	o11c	o12c	o13c	o14c	a1	r	o21c	o22c	a2	Y	S
5	0	-0.70404	1	0	-1	1		0		4	0
9	1	-1.01254	1	1	-1	1		0		4	0
23	0	0.14336	1	1	-1	1		1		5	0
24	1	1.07187	0	1	-1	1		0		1	0
27	0	-0.84515	0	0	-1	1		1		4	0
30	0	-0.21032	0	1	1	0	7	0	1	1	1

r : Response to Stage I treatment (1=Responders, 0=Non Responders)

O21c : Month during the school year at which the child showed inadequate response to the first-stage treatment, and hence entered the second-stage of the treatment.

O22c : Adherence to first-stage treatment(1=Adherence, 0=Non Adherence)

a2** :Second stage treatment

Y : School performance

S : Rerandomization indicators (1=Rerandomized, 0= Not Rerandomized)

**** A1=1=MED: Medication**
A1=-1=BMOD: Behavioral Intervention
A1=1=MED: Medication
A1=-1=BMOD: Behavioral Intervention

A2=1= INT: Intensify Medication
A2=1= intensify BMOD
A2=-1= Add BMOD
A2=-1=Add Medication

Q-Learning Regression Algorithm

i. Do a regression to learn how more deeply – tailor second stage treatment using O1 and O2(in ADHD SMART, this is only with non-responders).

ii. Assign each non responder the value $\hat{Y}_i$, an estimate of the outcome under the second line treatment A_{i^*} that yields best outcome for that person i. Responders get their observed Y_i.

iii. Using $\hat{Y}_i$ do a regression to learn how more deeply –tailor first stage treatment using O1

O1 : Other measures or Covariates of stage I
O2 : Other measures or Covariates of stage II

Step(i) of Q-Learning: Second Stage Treatment Tailoring?

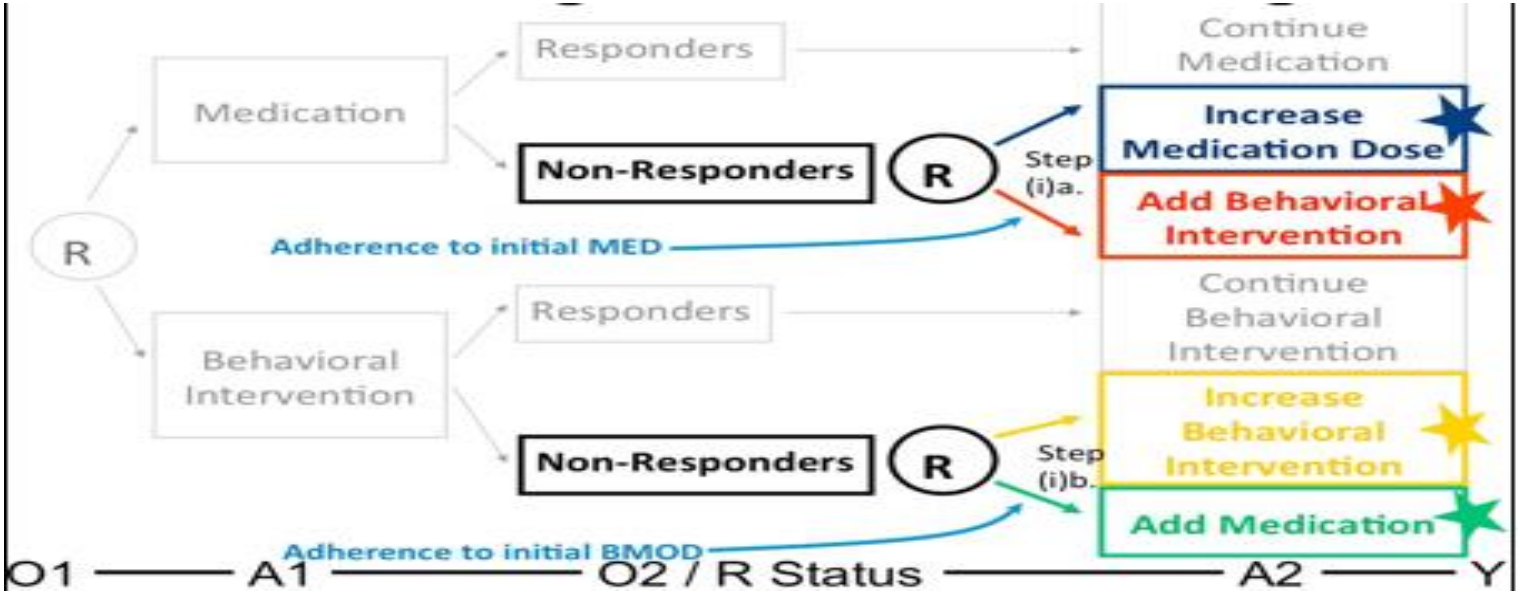

O1 : Other measures or Covariates of stage I

O2 : Other measures or Covariates of stage II

Step(i) of Q-Learning: Second Stage Treatment Tailoring?

– To accomplish a step (i) of Q learning, you may fit regression model such as this :

$$Y = b1 + b2\ o11c + b3\ o12c + b4\ o13c + b5\ o14c + b6\ o21c + b7\ a1 + b8\ o22 + \mathbf{b9\ a2*a1 + b10\ a2*o22 + b11\ a2} + error$$

From such a model we would learn if, for example o22(adherence to first stage txt) is a strong moderator of the effect of a2.

Note:

Y : School performance(Dependent Variable)

b1 : slope

b.. : regression Co- Efficients

Step (i) of Q-Learning

Learn best second-stage treatment for ***non-responders***

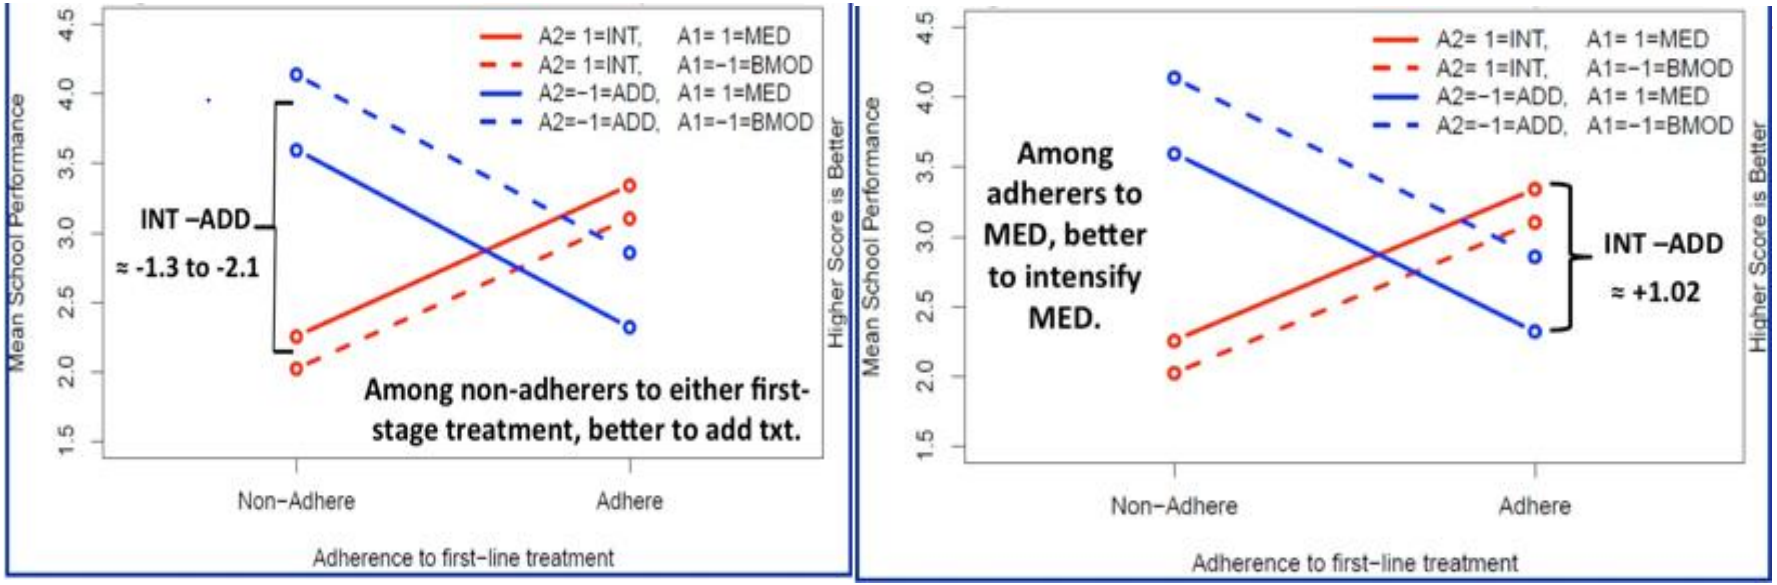

Note: **A1=1=MED: Medication** **A2=1= INT: Intensify Medication**

A1=-1=BMOD: Behavioral Intervention A2=1= intensify BMOD

A1=1=MED: Medication **A2=-1= Add BMOD**

A1=-1=BMOD: Behavioral Intervention A2=-1=Add Medication

Step (i) of Q-Learning

Learn best second-stage treatment for ***non-responders***

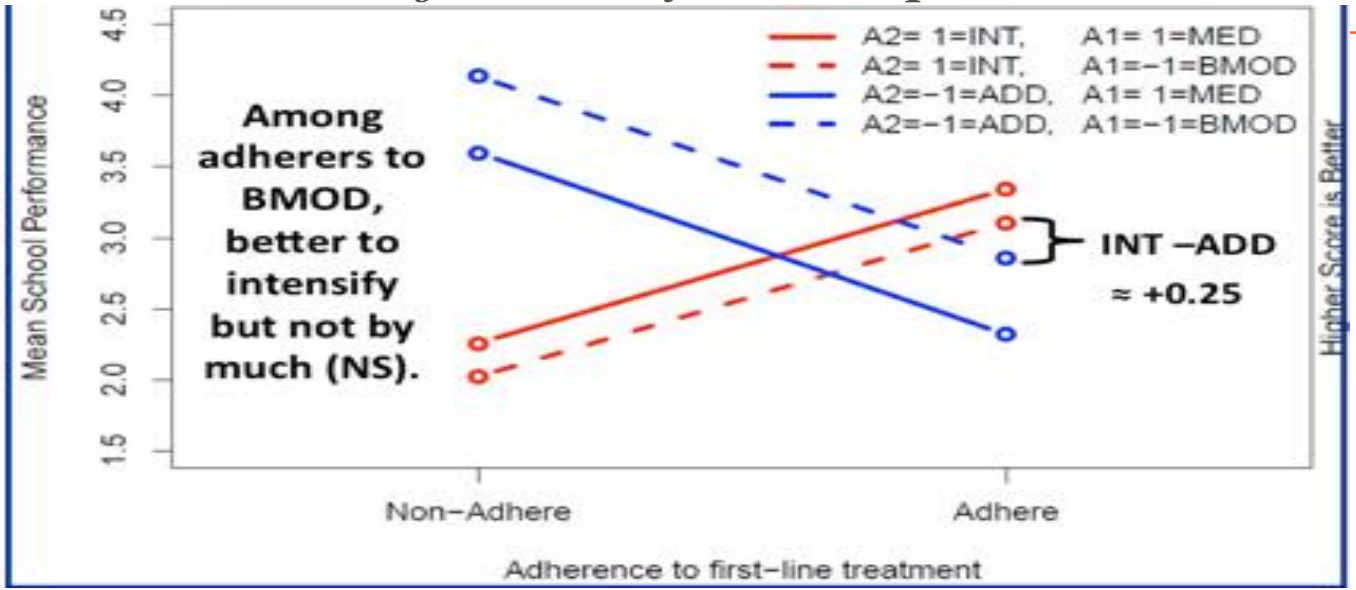

Note: **A1=1=MED: Medication** **A2=1= INT: Intensify Medication**

A1=-1=BMOD: Behavioral Intervention A2=1= intensify BMOD

A1=1=MED: Medication **A2=-1= Add BMOD**

A1=-1=BMOD: Behavioral Intervention A2=-1=Add Medication

Step (ii) of Q-Learning: Graphically

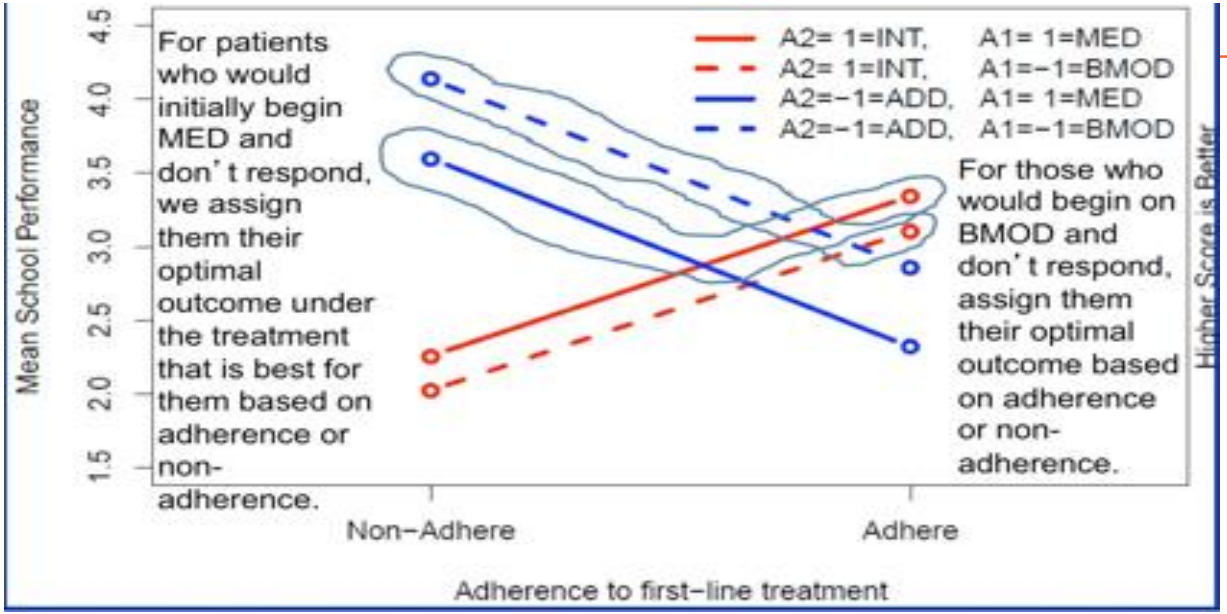

Note: **A1=1=MED: Medication** **A2=1= INT: Intensify Medication**

A1=-1=BMOD: Behavioral Intervention A2=1= intensify BMOD

A1=1=MED: Medication **A2=-1= Add BMOD**

A1=-1=BMOD: Behavioral Intervention A2=-1=Add Medication

Step (iii) of Q-Learning:

First-stage treatment tailoring?

– To accomplish step (iii) of Q learning you may fit a regression model such as this:

$$\hat{Y} = b1 + b2\ o11c + b3\ o12c + b4\ o14c + b5\ o13 + b6\ o13{*}a1 + b7\ a1 + error$$

– From such model we would learn if, for example, o13(prior MED) is a strongest moderator of the effect of a1 such that it would make a good tailoring variable.

Note:

Y : School performance(Dependent Variable)

b1: slope

b. : regression Co- Efficients

Step (iii) Q-Learning:

Learn optimal first treatment for all given optimal future txt

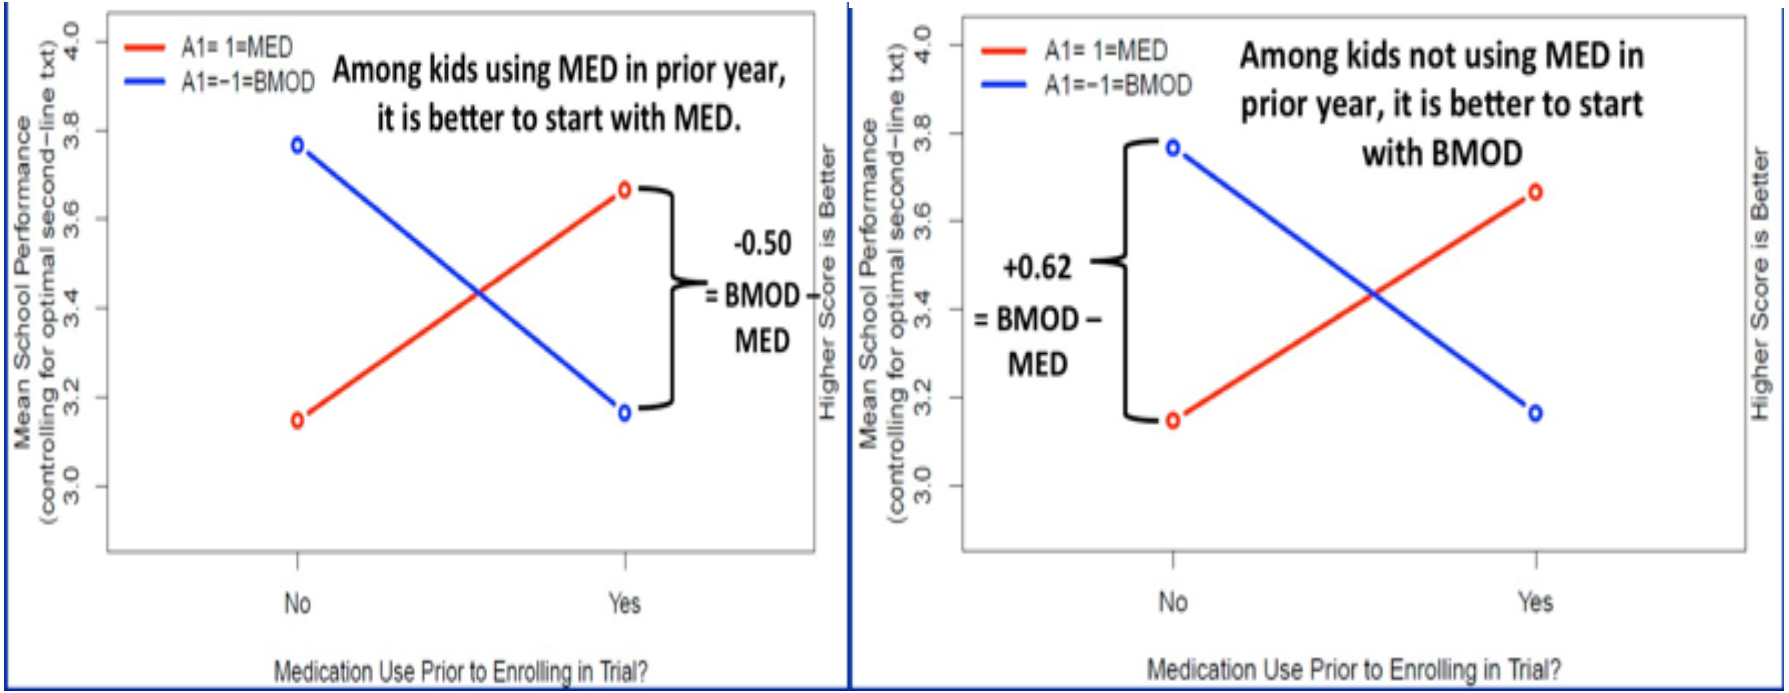

Note: **A1=1=MED: Medication** **A1=-1=BMOD: Behavioral Intervention**

PROC QLEARN

SAS add-on: PROC QLEARN

What do we provide PROC QLEARN?

1. Data set with O1, A1, R, O2, A2, Y
2. The second stage regression model
   - Y ~ O1, A1, O2, A2
   - Specify among who to do this regression (e.g., non-responders in ADHD SMART)
3. The first stage regression model
   - $\hat{Y}$ ~ O1, A1

SAS add-on: PROC QLEARN

What does PROC QLEARN do?

1. It implements Step (i): Second-stage Regression
   - Provides us with second-stage regression parameter estimates
2. It implements Step (ii): Obtains $\hat{Y}$
   - It assigns non-responders the outcome under the best second-stage treatment based on Step (i).
   - It assigns responders their observed outcome.
3. It implements Step (iii): First-stage Regression
   - Provides first-stage regression parameter estimates
   - Provides appropriate confidence intervals for the first-stage regression parameter estimates

PROC QLEARN

SAS add-on: PROC QLEARN

What does the syntax look like?

PROC QLEARN <options for input> ;
MAIN1 variables ;
TAILOR1 variables ;
MAIN2 variables ;
TAILOR2 variables ;
RESPONSE variable ;
STG1TRT variable ;
STG2TRT variable ;
STG2SAMPLE variable ;
ALPHA value ;
RUN;

```
data dat11; set dat1; S = 1-R; run;
proc qlearn data=dat11
   main1 o11c o12c o14c;
   tailor1 o13;
   main2  o11c o12c o13c o14c o21c;
   tailor2 a1 o22;
   stg2sample s;
   response y;
   stg1trt a1;
   stg2trt a2;
run;
```

contrasts1=contrasts1

deriveci;

This requests contrasts of interest (in GENMOD they are call "estimates". See next slide for details.

This requests confidence intervals by Laber and Murphy (2012)

This will ask SAS to fit the following two regressions:

Stage 2: $Y = b1 + b2\ o11c + b3\ o12c + b4\ o13c + b5\ o14c + b6\ o21c + b7\ a1 + b8\ o22 + b9\ a1*a2 + b10\ o22*a2 + b11\ a2 + error$ (fit only with non-responders, S=1)

Stage 1: $\hat{Y} = b1 + b2\ o11c + b3\ o12c + b4\ o14c + b5\ o13 + b6\ o13*a1 + b7\ a1 + error$

PROC QLEARN: Contrast Matrix

```
 data contrasts1;
Input M1 M2 M3 M4 M5 M6 M7;
datalines;
1 0 0 0 1  1  1    /*Mean outcome under BMOD for children with prior medication*/
1 0 0 0 1 -1 -1   /*Mean outcome under MED for children with prior medication*/
0 0 0 0 0  2  2   /*Mean Diff(BMOD - MED) for children with prior medication*/
1 0 0 0 0  0  1   /*Mean outcome under BMOD for children with no prior medication*/
1 0 0 0 0  0 -1   /*Mean outcome under MED for children with no prior medication*/
0 0 0 0 0  0  2   /*Mean Diff(BMOD - MED) for children with no prior medication*/
;
Run;
```

PROC QLEARN: Results

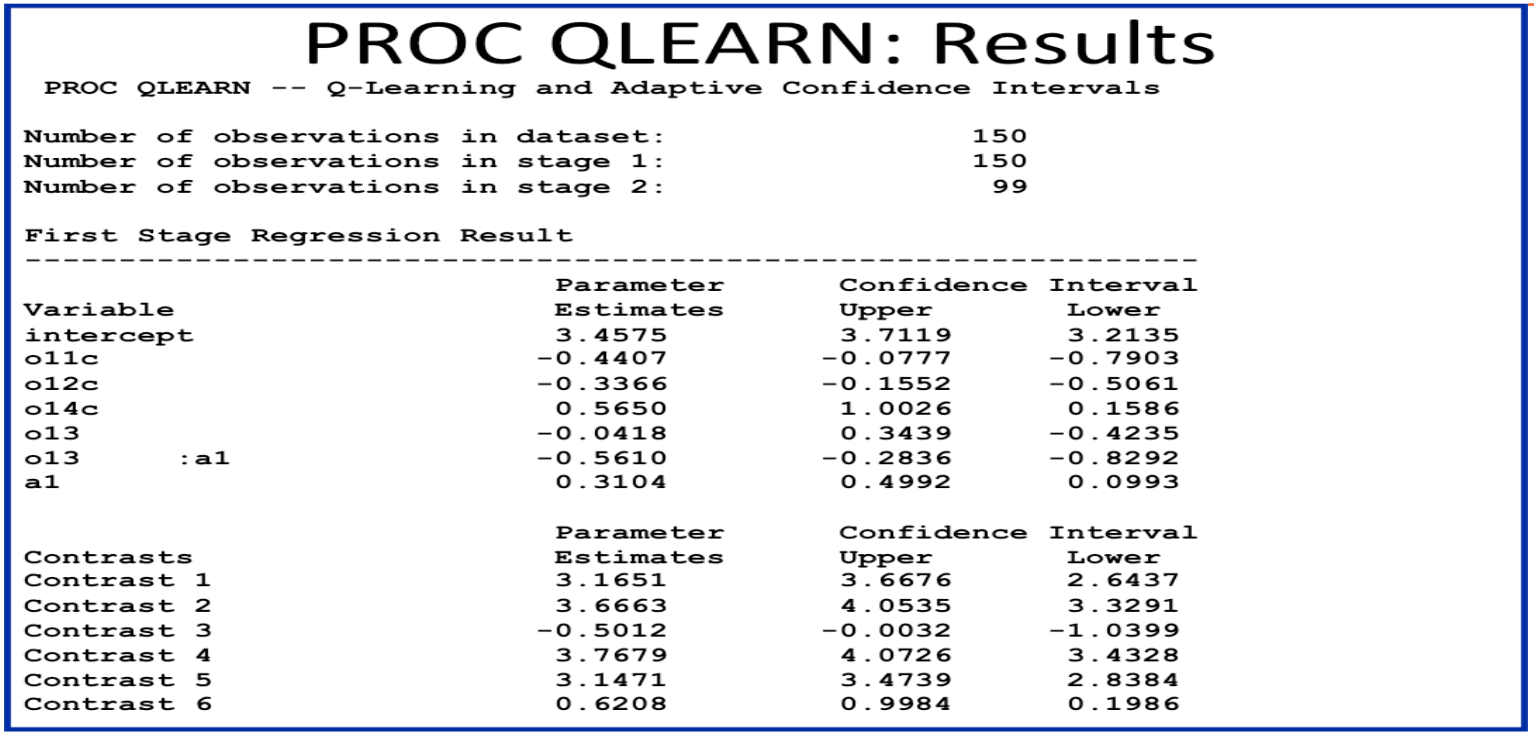

Interpretation

What did we learn with Q-learning?

Adaptive Intervention Proposal

- If the child used MED in prior year, then begin with MED; otherwise, begin with BMOD.
- If the child is non-responsive and non-adherent to either first-line treatment, then AUGMENT with the other treatment option.
- If the child is non-responsive but adherent to either first-line treatment, then it is better to INTENSIFY first-line treatment.
- If the child is responsive to first-line treatment, then CONTINUE first-line treatment.

This Q-learning analysis was done with simulated/altered data.

References used

http://www-personal.umich.edu/~dalmiral/slides/SREE_Getting-SMART_Syllabus-and-Slides.pdf

http://www2.cscc.unc.edu/impact7/system/files/Workshop%20on%20Personalized%20Medicine%20and%20Dynamic%20Treatment%20Regimes_0.pdf

http://docplayer.net/41234675-Sas-proc-qlearn-users-guide-version-1-0-0.html

http://www-personal.umich.edu/~dalmiral/software/mw_workshop_files/SAS%20Code/adhd_simulated_data.txt

Thank You