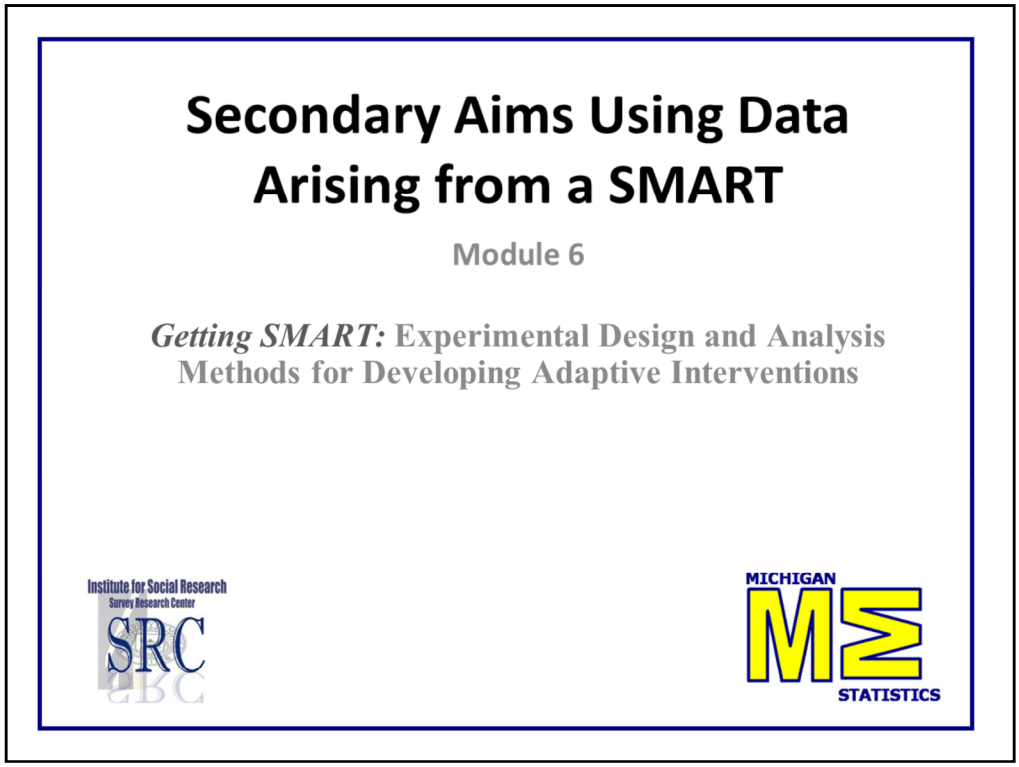

This is the 6th (final) module of a 2-day 6-module workshop on experimental designs for building optimal adaptive interventions.

By now, you know what an AI is. You have discussed why they are important in terms of managing chronic disorders (indeed, an AI formalizes the type of clinical practice taking place today). You have been introduced to the SMART clinical trial design, the rationale for SMARTs, and some important SMART design principles. Also, you have been introduced to typical primary aims and their associated data analysis methods.

In this module, we are going discuss Q-learning, a new type of data analysis method used as a secondary research aims using data arising from SMART studies.

Outline

- Discuss what is a "more deeply-tailored AI"
- Review of auxiliary data typical of SMARTs
- Q-Learning

Outline

- Discuss what is a "more deeply-tailored AI"
- Review of auxiliary data typical of SMARTs
- Q-Learning

What is a More Deeply-Tailored AI?

- To understand this, we first review what are the "embedded AIs" within the ADHD SMART study
- Recall there are 4 SMART-design embedded AIs.

Remember ADHD SMART?

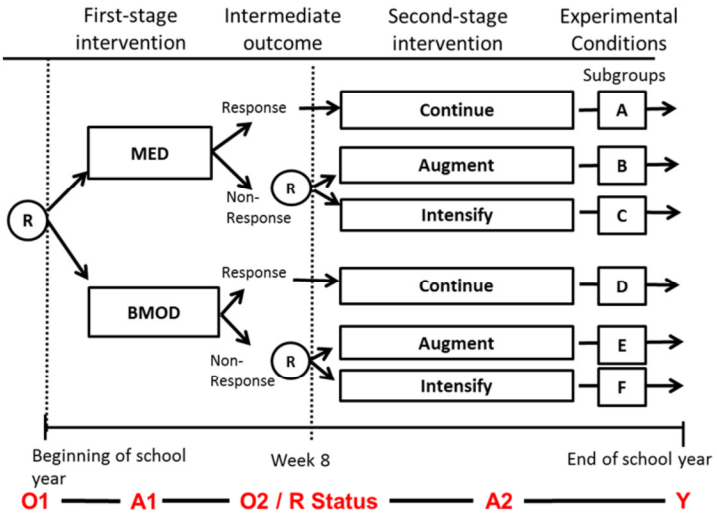

4 Embedded AIs

Here is AI #1

At the beginning of school year

Stage 1 = {**MED**},

Then, every month,

starting at week 8

IF response status = {NR}

THEN Stage 2 = {**AUGMENT**}

ELSE IF response status = {R}

THEN **CONTINUE** Stage 1

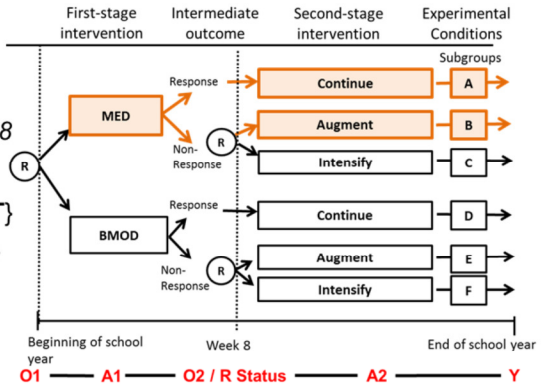

4 Embedded AIs

Here is AI #2

At the beginning of school year

Stage 1 = {**BMOD**},

Then, every month,

starting at week 8

IF response status = {NR}

THEN Stage 2 = {**INTENSIFY**}

ELSE IF response status = {R}

THEN **CONTINUE** Stage 1

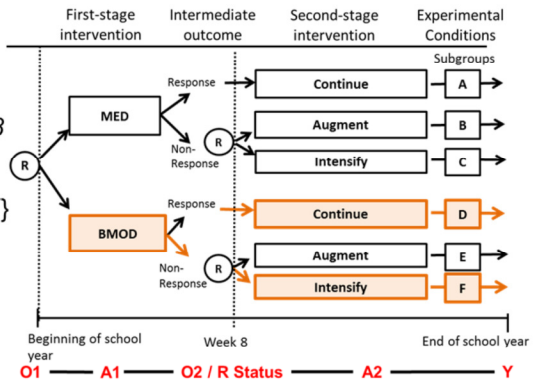

4 Embedded AIs

Here is AI #3

At the beginning of school year

Stage 1 = {**MED**},

Then, every month,

starting at week 8

IF response status = {NR}

THEN Stage 2 = {**INTENSIFY**}

ELSE IF response status = {R}

THEN **CONTINUE** Stage 1

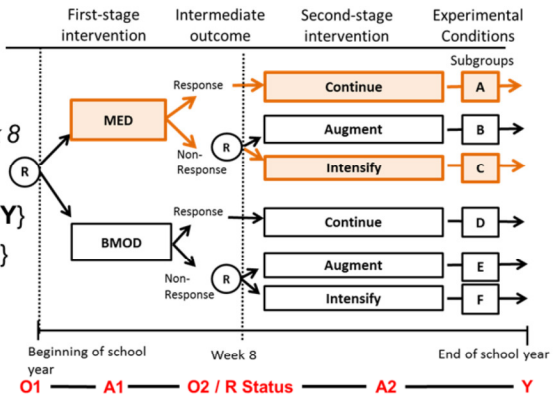

4 Embedded AIs

Here is AI #4

At the beginning of school year

Stage 1 = {**BMOD**},

Then, every month,

starting at week 8

IF response status = {NR}

THEN Stage 2 = {**AUGMENT**}

ELSE IF response status = {R}

THEN **CONTINUE** Stage 1

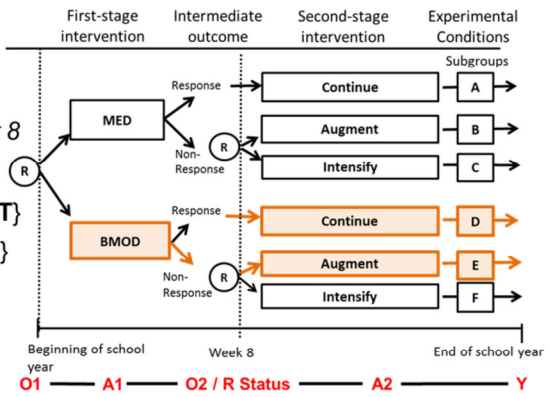

What is a More Deeply-Tailored AI?

- There is only 1 tailoring variable embedded by design in the ADHD SMART
  - Response Status
- A more deeply tailored AI is a sequence of decision rules that include tailoring variables **beyond those** embedded in the SMART by design.
  - In our case, this would be an AI that includes tailoring variables beyond response-status.

Why Consider More Deeply-Tailored AIs?

1. It may be that some participants may benefit more from starting on MED vs starting on BMOD.
   - For example: those who have used MED in the past
2. Certain types of non-responders may benefit more from AUGMENT vs. INTENSIFY
   - For example, those who do not adhere to initial treatment

A More Deeply-Tailored AI
Might Look Like This:

At the beginning of school year

IF medication in prior year = {NO}
THEN stage 1 = {BMOD}.
ELSE IF medication in prior year ={YES}
THEN stage 1 = {MED}

Then, every month,
beginning at week 8

IF response status to Stage 1 = {NR}
THEN IF adherence to stage 1 = {NO},
THEN Stage 2 = {AUGMENT}.
ELSE Stage 2 = {AUGMENT} or {ENHANCE}.
ELSE CONTINUE Stage 1.

A More Deeply-Tailored AI
Might Look Like This:

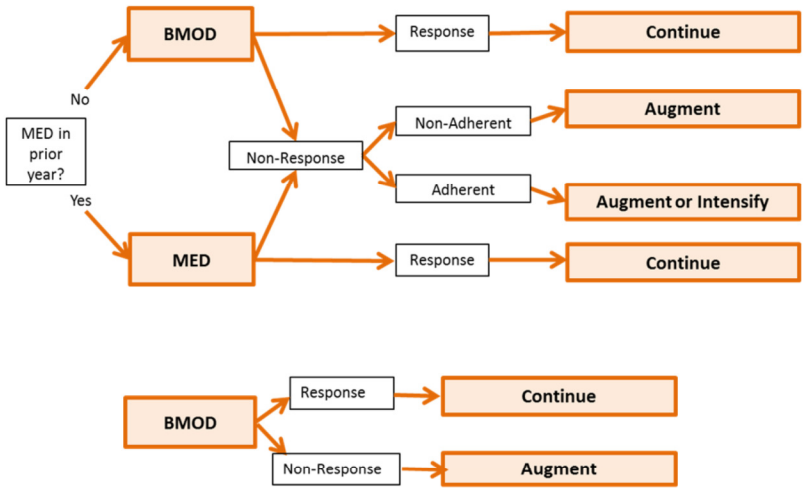

Outline

- Discuss what is a "more deeply-tailored AI"
- Review of auxiliary data typical of SMARTs
- Q-Learning

The remaining slides in this Module are devoted to understanding how to use auxiliary data arising from a SMART with a regression method known as Q-learning to develop/learn/discover a more deeply-tailored AIs such as the one shown on the previous slide.

But first, what do we mean by auxiliary data?

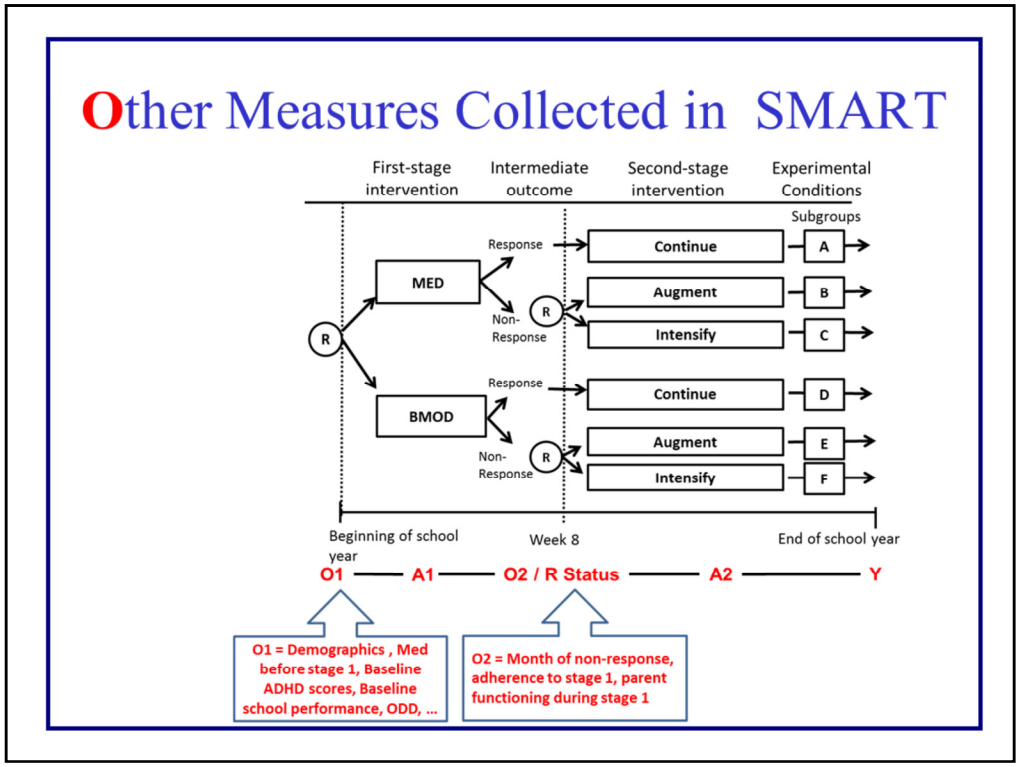

In addition to standard outcomes scales/measures, many other things could me measured during initial treatment (in this SMART study) that could be used in secondary analyses to more deeply tailor/individualize subsequent treatment, including:

Allegiance/rapport of individual with the psychologist/psychiatrist,
Environmental outcomes (parent outcomes, …),
Ecological momentary assessments (daily/weekly substance use patterns, rituals, etc.)

Notice that some of the O2 measures may be available for non-responders, but not available (e.g., “structurally missing”) for responders: an example of this in this ADHD study is the time until non-response!

How to Use O1 and O2?

- We can use the auxiliary data O1 to help decide who would benefit more from MED vs. BMOD.
- We can use the auxiliary data O1 and O2 to help decide who (among the non-responders) would benefit more from INTENSIFY vs. AUGMENT.

Outline

- Discuss what is a "more deeply-tailored AI"
- Review of auxiliary data typical of SMARTs
- Q-Learning

The remaining slides in this Module are devoted to understanding how to use auxiliary data arising from a SMART with a regression method known as Q-learning to develop/learn/discover a more deeply-tailored ATS such as the one shown on the previous slide.

Q-Learning

- Q-Learning is an extension of regression to sequential treatments.
- Q is for learning about the "Quality" of the AI.
- Q-Learning results in a ***proposal*** for an AI with greater individualization.
  - Namely, one that includes more tailoring variables than those AIs embedded in the SMART by design
- You can use a subsequent trial (i.e., RCT) to evaluate the proposed AI versus a suitable control (e.g., usual care).

This is an idea borrowed from computer scientists.

3-Steps in Q-Learning

Step 1: Regression 1

Are O1, A1, and O2 useful in making decisions about second-stage tactics for NR?

Whether O1, A1, and O2 are useful in deciding who would benefit from Augment vs. Intensify.

Step 2: Calculate $\hat{Y}i$

What would be the expected outcome for each non-responder if he/she had received the best tactic given his/her O1, A1, and O2?

$\hat{Y}i$ is the estimated optimal outcome under the best second-stage option for non-responders.

For responders $\hat{Y}i = Yi$

Step 3: Regression 2

Are O1 useful in making decisions about the initial treatment, assuming the in the future the best tactic will be used for non-responders?

Use $\hat{Y}i$ form Step 2 as the outcome.

Whether O1 is useful in deciding who would benefit from BMOD vs. MED

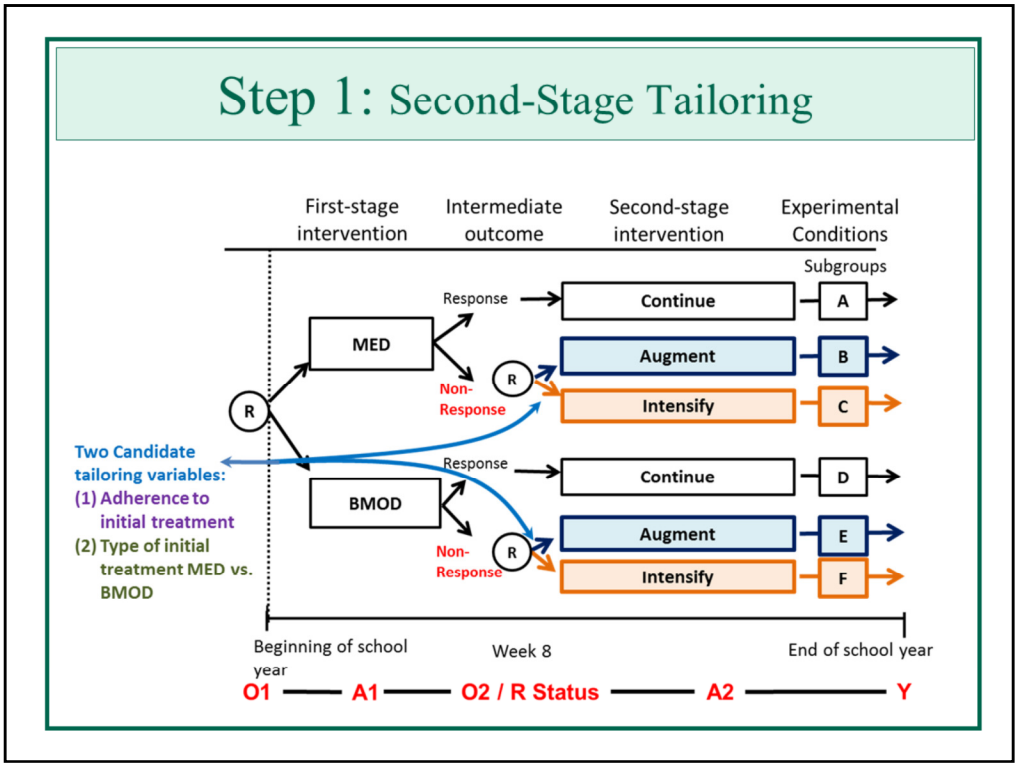

In this step we focus on adherence to first stage as a tailoring variable for the second-stage tactic offered to non-responders.
We also focus on the type of first stage treatment as a tailoring variables

Here we want to know whether adherence to the initial treatment is useful for deciding whether to augment or intensify the initial treatment.
In other words, we want to know if the second-stage tactic should be tailored based on the level of adherence to the first stage.

Step 1: Second-Stage Tailoring

- In this step, we seek to address 2 questions:

1. Can we use information about **adherence** to initial treatment to select a tactic for non-responders?
2. Can we use information about the **initial treatment** to select a tactic for non-responders?

- Fit a moderated regression model using data from non-responders to address these questions:

$$Y = b1 + b2\ O11c + b3\ O12c + b4\ O13c + b5\ O14c + b6\ O21c + b7\ A1 + b8\ O22 + \mathbf{b9\ A2} + \mathbf{b10\ A2{*}A1} + \mathbf{b11\ A2{*}O22} + \text{error}$$

See next slide for more details …

Step 1: Second-Stage Tailoring

$$Y = b1 + b2\ O11c + b3\ O12c + b4\ O13c + b5\ O14c + b6\ O21c + b7\ A1 + b8\ O22 + \mathbf{b9\ A2} + \mathbf{b10\ A2{*}A1} + \mathbf{b11\ A2{*}O22} + \text{error}$$

Baseline covariates

Intermediate covariates

Why interactions?

Because we want to know whether and how the two candidate tailoring variables moderate the effect of A2.

This model will help us

(a) Determine whether the best second stage tactics varies depending on the tailoring variables

(b) Identify the best tactic depending on the level of the tailoring variable

A1 = stage 1 options: -1=MED; 1=BMOD
A2 = stage 2 tactic: -1=ADD; 1=INTSFY
O11 = baseline ODD: 1=yes; 0=no
O12 = baseline ADHD score: hi is better
O13 = med before stage 1: 1=yes; 0=no
O14 = race: 1=white; 0=non-white
O21 = # of months until non-response:
O22 = adherence to stage 1: 1=yes; 0=no
Y = end of year school performance

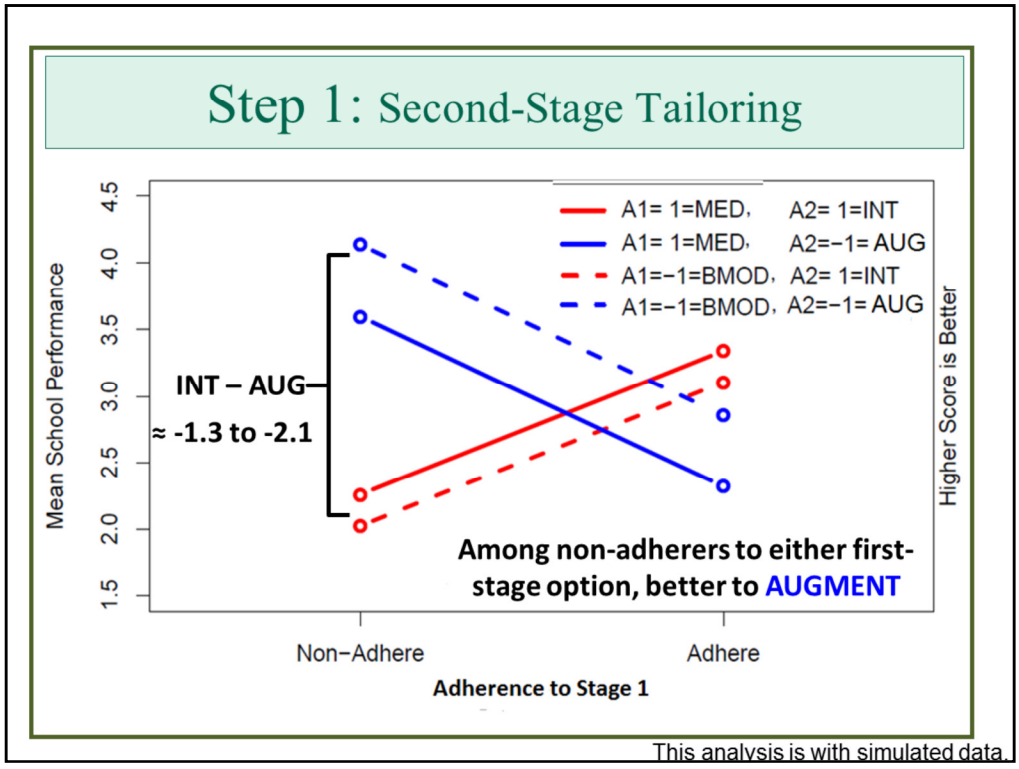

This is what we might learn from a regression such as the one shown on the previous slide.

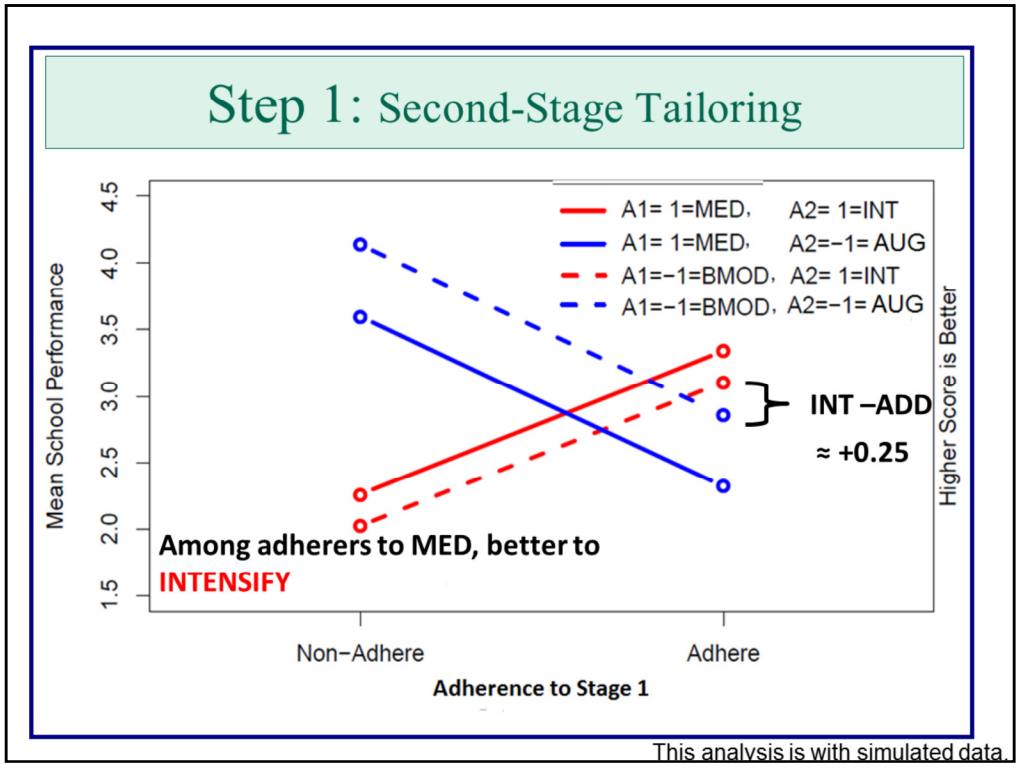

This is what we might learn from a regression such as the one shown on the previous slide.

Step 2: Predicted Outcome Under the Best Stage 2 Option

- In this step, we assign each non-responder the value $\hat{Y}_i$
  - What would be the outcome if a non-responder received the best second-stage tactic given his/her initial treatment and adherence?
    - Based on Regression 1, we identified the best stage 2 tactic for any given level of the tailoring variables.
    - We use these results to estimate what would be the outcome if we had given each non-responder the best stage 2 tactic,
      - » Given his/her observed values on the tailoring variables
    - This estimate would be $\hat{Y}_i$
    - Responders get their observed Y_i.

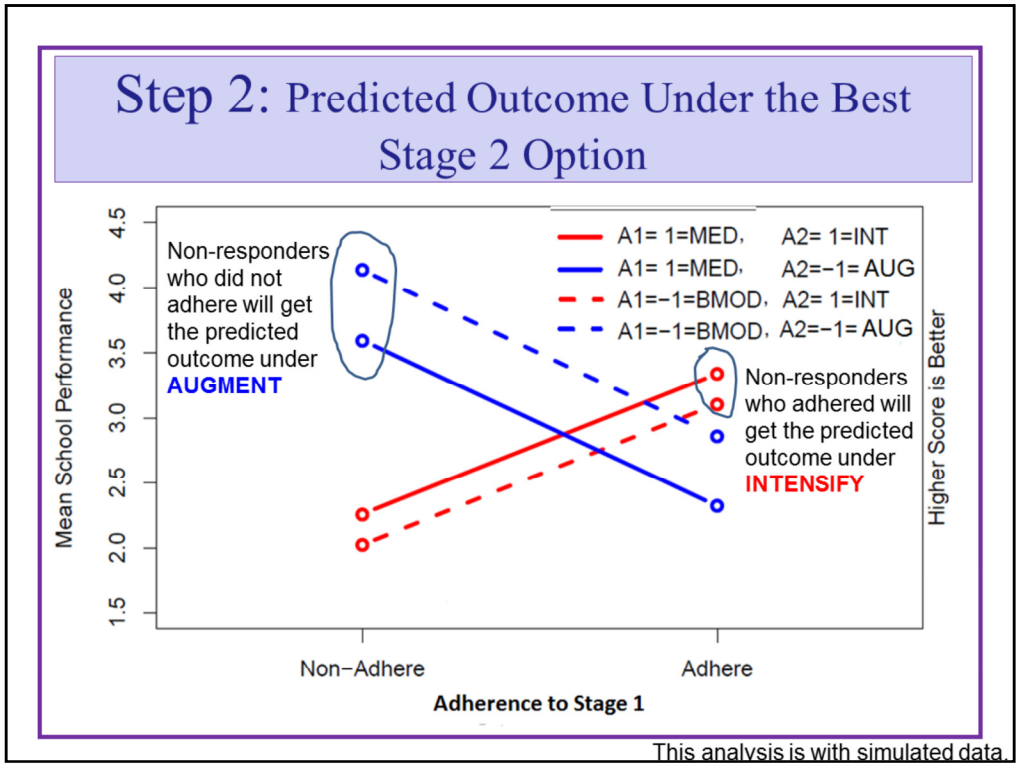

This mean that for non-responders who adhered we give them the predicted outcome that they would get if they had been assigned to Augment
And, for non-responders who adhered, we give them the predicted outcome that they would get if they had been assigned to intensify.

Step 3: Move Backwards to First-Stage Tailoring

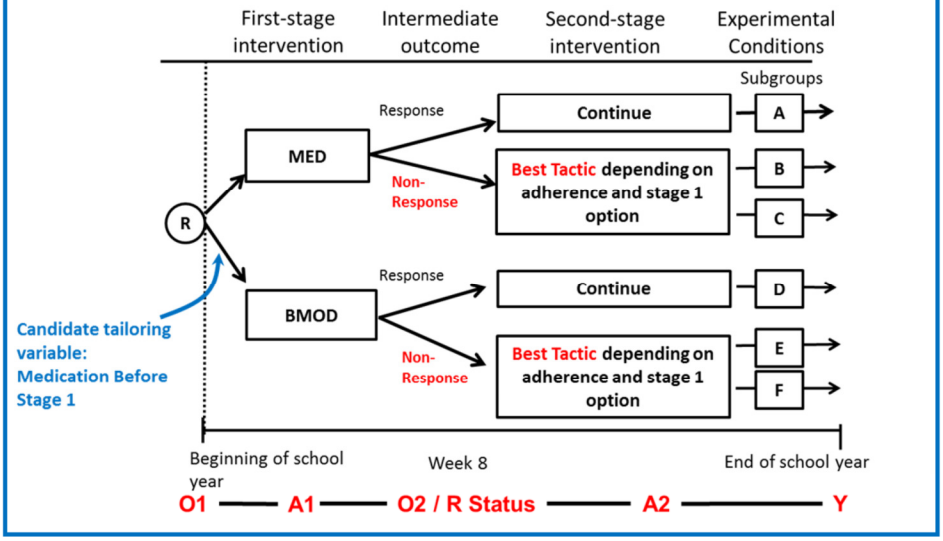

Step 3: Move Backwards to First-Stage Tailoring

- In this step, we seek to address the following question:
  Can we use information about **medication in prior year (O13)** to select a first-stage option?
- Assuming that in the future, non-responders get the best subsequent tactic.
- This is done by using $\hat{Y}_i$ as the outcome in the regression where we explore the usefulness of **O13** for making decisions about first-stage options.

Step 3: Move Backwards to First-Stage Tailoring

Fit the following regression model:

$$\hat{Y} = b1 + b2\ O11c + b3\ O12c + b4\ O14c + b5\ \mathbf{O13} + b6\ \mathbf{A1} + b7\ \mathbf{O13}*\mathbf{A1} + \text{error}$$

Controlling for stage 2 tactic

Why interaction?

Because we want to know whether and how **medication in prior year (O13)** moderates the effect of **stage 1 intervention options** (**A1**).

This model will help us

(a) Determine whether the best first stage option varies depending on whether or not the child received medication in prior year

(b) Identify the best first stage option for children who received med in prior year vs. those who did not.

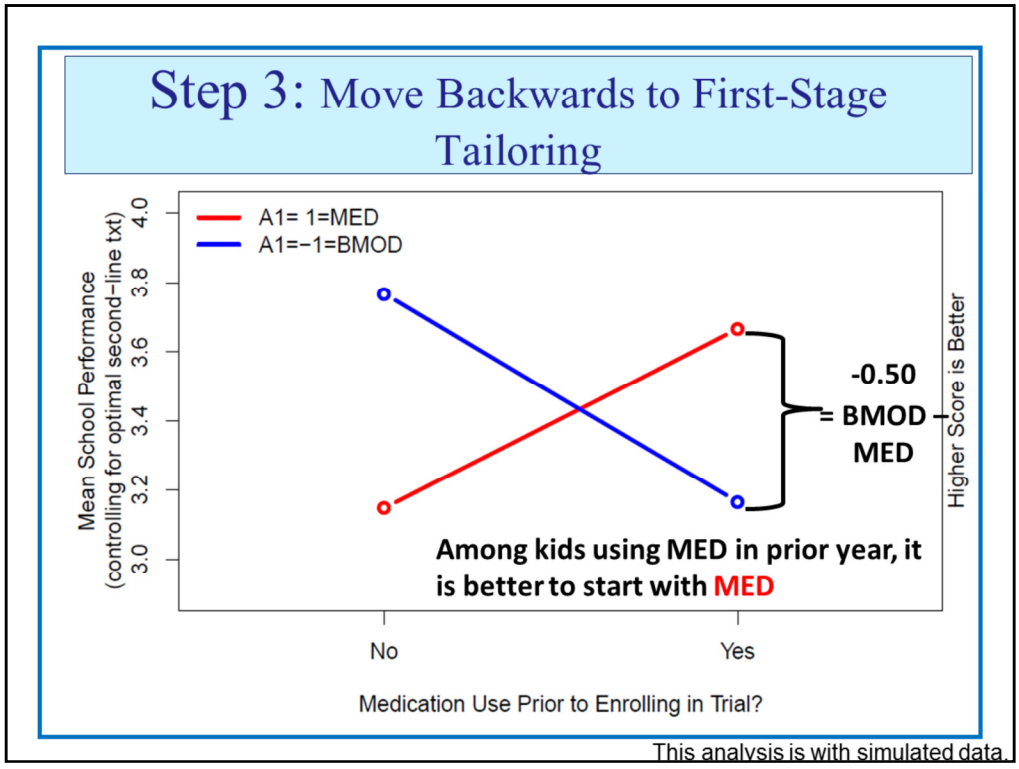

So, we should assign MED to kids with MED in prior year

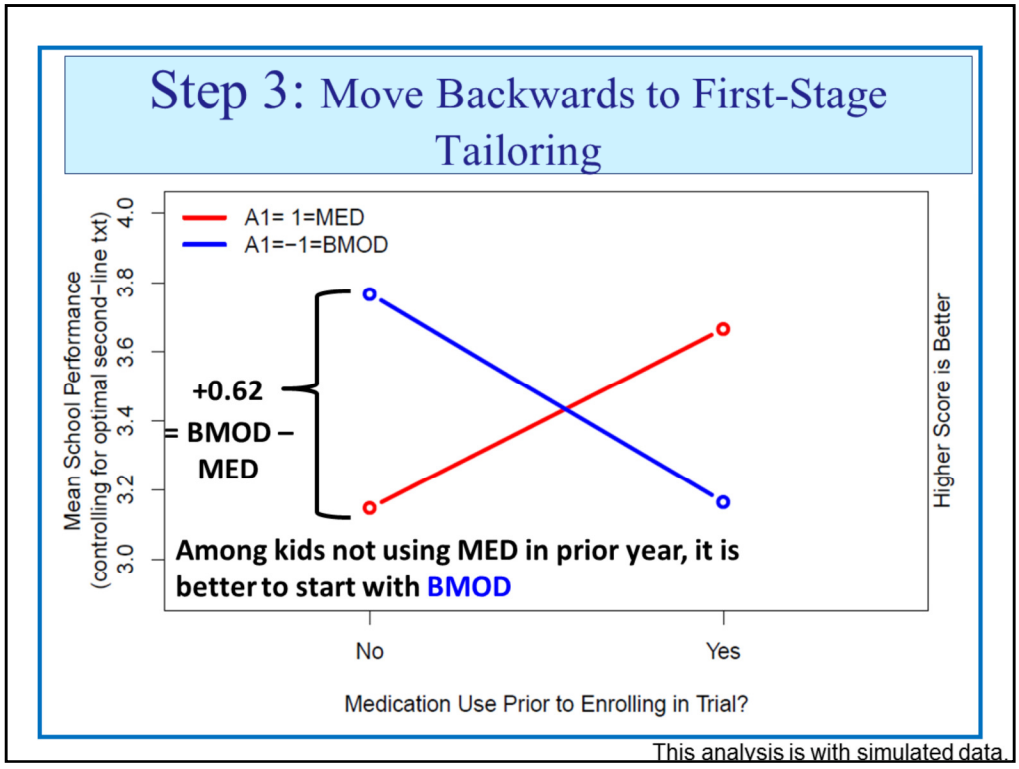

And, we should assign BMOD to kids who did not have MED in the prior year

3-Steps in Q-Learning

Summary of Results

Step 1: Regression 1

A1 & O22 are useful in making decisions about second-stage tactics for NR!

NR Non-Adherent → Augment

NR adherent to MED→ Intensify

NR adherent to BMOD→ Augment/Intensify

Step 2: Calculate $\hat{Y}i$

We calculated the expected outcome for each non-responder if he/she had received the best tactic given his/her A1, and O2.

NR Non-Adherent → Predicted outcome under Augment

NR adherent → Predicted outcome under Intensify

Step 3: Regression 2

O13 is useful in making decisions about the initial treatment, assuming the in the future the best tactic will be used for non-responders.

YES MED→ start with MED

NO MED→ start with BMOD

The Estimated More Deeply-Tailored AI is:

At the beginning of school year

IF medication in prior year = {NO}

THEN stage 1 = {BMOD}.

ELSE IF medication in prior year ={YES}

THEN stage 1 = {MED}

Then, every month,

beginning at week 8

• • •

The Estimated More Deeply-Tailored AI is:

… *Then, every month,*
beginning at week 8

IF response status to Stage 1 = {NR}
Then,
IF adherence to MED or BMOD= {NO},
THEN Stage 2 = {AUGMENT}.
Else IF adherence to MED = {YES},
THEN Stage 2 = {INTENSIFY}.
Else IF adherence to BEMOD = {YES},
THEN Stage 2 = {AUGMENT} or {ENHANCE}.
ELSE IF **response status** to Stage 1 = {R}
Then, CONTINUE Stage 1.

SAS Software to Use Q-Learning

- We next show you how to do
  - Step 1 using regression
  - Steps 2 and 3 using a SAS add-on known as PROC QLEARN
- We will use the two example regression models shown previously.

SAS Software to Use Q-Learning

Step 1 using regression

$$Y = b1 + b2\ O11c + b3\ O12c + b4\ O13c + b5\ O14c + b6\ O21c + b7\ A1 + b8\ O22 + \mathbf{b9\ A2} + \mathbf{b10\ A2{*}A1} + \mathbf{b11\ A2{*}O22} + \text{error}$$

```
* use only non-responders;
data dat10; set dat1; if R=0; run;

proc genmod data = dat10;
  model y = o11c o12c o13c o14c o21c a1 o22 a2 a2*a1 a2*o22;

  * diff INTENSIFY vs. ADD when stage 1 = MED by ADH status;
  estimate 'INT vs ADD for NR ADH MED'      a2 2 a2*a1 -2 a2*o22 2 ;
  estimate 'INT vs ADD for NR Non-ADH MED'  a2 2 a2*a1 -2 a2*o22 0 ;
  * diff INTENSIFY vs. ADD when stage 1 = BMOD by ADH status;
  estimate 'INT vs ADD for NR ADH BMOD'     a2 2 a2*a1 2 a2*o22  2 ;
  estimate 'INT vs ADD for NR Non-ADH BMOD' a2 2 a2*a1 2 a2*o22  0 ;
run;
```

SAS Software to Use Q-Learning

Step 1 using regression

$$Y = b1 + b2\ O11c + b3\ O12c + b4\ O13c + b5\ O14c + b6\ O21c + b7\ A1 + b8\ O22 + \mathbf{b9\ A2} + \mathbf{b10\ A2{*}A1} + \mathbf{b11\ A2{*}O22} + \text{error}$$

Contrast Estimate Results

Label	Estimate	95% Conf Limits Lower	Upper	P-value
INT vs ADD for NR ADH MED	1.0240	0.4131	1.6350	<.0001
INT vs ADD for NR Non-ADH MED	-1.3412	-1.9896	-0.6927	<.0001
INT vs ADD for NR ADH BMOD	0.2503	-0.3950	0.8956	0.4471
INT vs ADD for NR Non-ADH BMOD	-2.1149	-2.7050	-1.5248	<.0001

Contrast 1: Among non-responders who adhere to MED, it is better to INTENSIFY treatment rather than ADD a different treatment (positive effect).

Contrast 2: Among non-responders who do not adhere to MED, it is better to ADD than to INTENSIFY (negative effect).

Contrast 3: Among non-responders who adhere to BMOD, it is better to INTENSIFY than ADD (but this effect is not significant at 0.05; p-value=0.45).

Contrast 4: Among non-responders who do not adhere to BMOD, it is better to ADD than to intensify (negative effect).

Try it yourself in SAS

- Go to the file:

  sas_code_modules_4_5_and_6_ADHD.doc
- Copy the SAS code on **Page 11**
- Paste into SAS Enhanced Editor window
- Press F8 or click the Submit (little running guy)

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 sub-sample for this regression (e.g., non-responders in ADHD SMART)

3. The first stage regression model

•$\hat{Y}$ ~ O1, A1

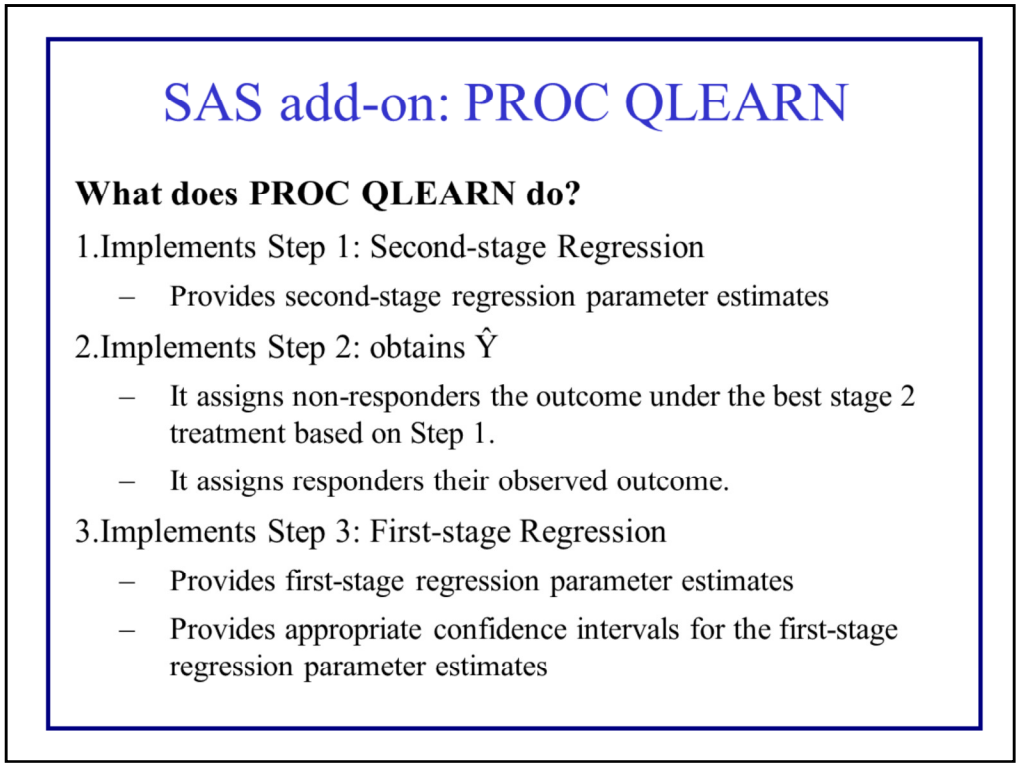

In Step1 PROC QLEARN just reproduces the second stage regression you did by hand.

Step 3 implements a special bootstrap procedure (Laber and Murphy, 2012; JASA) to produce appropriate statistical inferences (confidence intervals) concerning the first-stage regression parameters.

Laber and Murphy (2012; JASA) call these “adaptive confidence intervals”.

SAS add-on: PROC QLEARN

Model Specification

```
PROC QLEARN <options for input> ;
   MAIN1 variables;
   TAILOR1 variables;
   MAIN2 variables;
   TAILOR2 variables;
   RESPONSE variable;
   STG1TRT variable;      *Must be coded -1/+1
   STG2TRT variable;      *Must be coded -1/+1
   STG2SAMPLE variable;   *0/1 indicator specifying sample used for stage 2
   ALPHA value;           *Type-I error to calculate CI for stage 1 reg.
RUN;

STAGE2 REGRESSION RESULTS ARE INPUT LIKE THIS:
INT + MAIN2 + TAILOR2 + TAILOR2*STG2TRT + STG2TRT

STAGE1 REGRESSION RESULTS ARE INTPUT LIKE THIS:
INT + MAIN1 + TAILOR1 + TAILOR1*STG1TRT + STG1TRT
```

MAIN1 and TAILOR1 are used to specify the first-stage regression (next slide explains)
MAIN2 and TAILOR2 are used to specify the second-stage regression (next slide explains)
RESPONSE specifies the outcome variable Y
STG1TRT gives SAS the name of the A1 variable, must be coded -1/+1
STG2TRT gives SAS the name of the A2 variable, must be coded -1/+1
STG2SAMPLE is a 0/1 indicator variable specifying (if equal to 1) the sample used for the second-stage regression
ALPHA specifies the Type-I error used to calculate the confidence intervals for the first-stage regression

SAS add-on: PROC QLEARN

Model Specification:

•Recall our stage 2 (Step 1) model for non-responders

```
data dat10; set dat1; if R=0; run; * use only NR's;
proc genmod data = dat10;
    model y = o11c o12c o13c o14c o21c a1 o22 a2 a2*a1 a2*o22;
run;
```

•This is how this model will be specified for QLEARN

```
data dat11; set dat1; S = 1-R; run; * use only NR's;
proc qlearn data=dat11;
   ...[next slide]
   main2   o11c o12c o13c o14c o21c;
   tailor2 a1 o22;
   stg2sample s;
   response y;
   stg2trt a2; ...[next slide]
run;
```

Notice this tells PROC QLEARN to do the second stage regression with only the S=1 participants which are the R=0 participants, which are the non-responders who were re-randomized.

In the next slide we show how to specify the first-stage regression by showing a complete specification for PROC QLEARN.

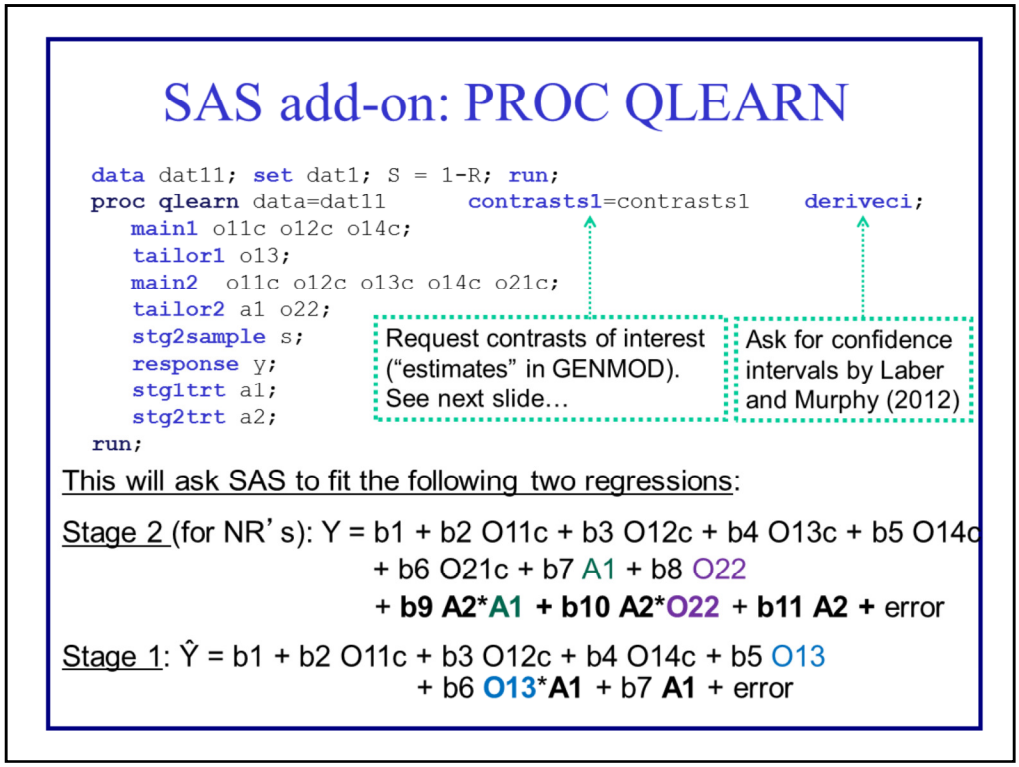

There are other options (optional) that we do not describe in the slide.
The User's Guide explains these in more detail.

SAS add-on: PROC QLEARN

How to Specify Contrast Matrix

```
data contrasts1;
  input M1 M2 M3 M4 M5 M6 M7;
  * cols correspond to the parameters in stage 1 model:
;
  * b1 + b2 O11c + b3 O12c + b4 O14c + b5 O13 + b6 A1*O13 + b7 A1;
  * each row corresponds to a different linear comb of the b's.;
  datalines;
  1 0 0 0 1  1  1  /*mean under BMOD for kids w prior med      */
  1 0 0 0 1 -1 -1  /*mean under MED for kids w prior med       */
  0 0 0 0 0  2  2  /*mean diff (BMOD-MED) kids w prior med     */
  1 0 0 0 0  0  1  /*mean under BMOD for kids w no prior med   */
  1 0 0 0 0  0 -1  /*mean under MED for kids w no prior med    */
  0 0 0 0 0  0  2  /*mean diff (BMOD - MED) kids w no prior med*/
 run;


estimate 'Mean und BMOD prior med' intercept 1 o13 1 a1*o13 1 a1 1;
```

** # cols = # of stage 1 parameters;
** some linear combos can be used to obtain mean outcomes for different children under initial BMOD vs initial MED.
whereas other linear combos can be used to
compare mean outcomes between MED vs BMOD for different children.

Try it yourself in SAS

- Go to the file:

  sas_code_modules_4_5_and_6_ADHD.doc

- Copy the SAS code on Page 13
  - This code defines the contrast matrix and runs PROC QLEARN
- Paste into SAS Enhanced Editor window
- Press F8 or click the Submit (little running guy)

SAS add-on: PROC QLEARN

First Stage Regression Result

Variable	Parameter Estimates	Confidence Interval Upper	Lower
intercept	3.4575	3.7119	3.2135
o11c	-0.4407	-0.0777	-0.7903
o12c	-0.3366	-0.1552	-0.5061
o14c	0.5650	1.0026	0.1586
o13	-0.0418	0.3439	-0.4235
o13 :a1	-0.5610	-0.2836	-0.8292
a1	0.3104	0.4992	0.0993

Contrasts	Parameter Estimates	Confidence Interval Upper	Lower
Contrast 1	3.1651	3.6676	2.6437
Contrast 2	3.6663	4.0535	3.3291
Contrast 3	-0.5012	-0.0032	-1.0399
Contrast 4	3.7679	4.0726	3.4328
Contrast 5	3.1471	3.4739	2.8384
Contrast 6	0.6208	0.9984	0.1986

My results for the estimates will be identical to yours. My results for the confidence intervals will be different from yours. This is because the confidence intervals by Laber and Murphy (2012) are based on a bootstrapping procedure that re-samples the data. Not shown on this slide (but it will show on your output screen are the parameter estimates for the stage 2 model). PROC QLEARN provides these so you can be sure you implemented the correct stage 2 model that you implemented previously. You should check you got the same answers as for the model you ran on Slide 20.

Interpretations:
You can see that medication in the prior year is, indeed, a significant tailoring variable. That is, the sign for BMOD vs MED changes (contrast 3 vs contrast 6) depending on the level of medication in the prior year. Since contrast 3 is borderline, it would not be surprising if some of you see that this interval covers zero.

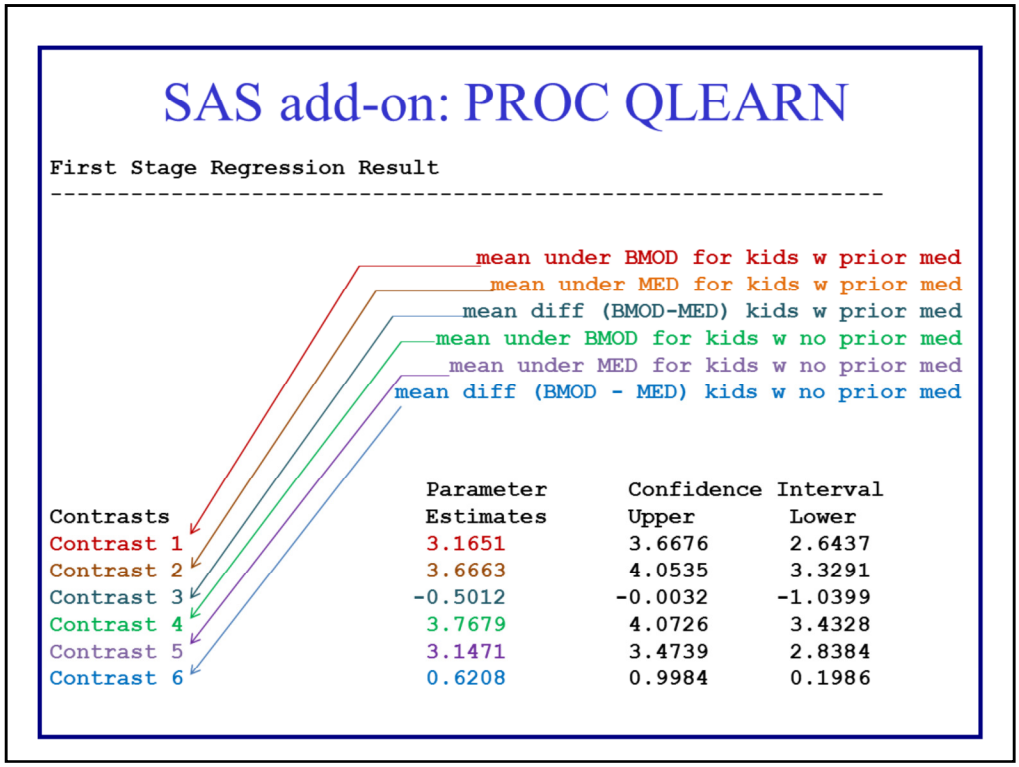

My results for the estimates will be identical to yours. My results for the confidence intervals will be different from yours. This is because the confidence intervals by Laber and Murphy (2012) are based on a bootstrapping procedure that re-samples the data. Not shown on this slide (but it will show on your output screen are the parameter estimates for the stage 2 model). PROC QLEARN provides these so you can be sure you implemented the correct stage 2 model that you implemented previously. You should check you got the same answers as for the model you ran on Slide 20.

Interpretations:
You can see that medication in the prior year is, indeed, a significant tailoring variable. That is, the sign for BMOD vs MED changes (contrast 3 vs contrast 6) depending on the level of medication in the prior year. Since contrast 3 is borderline, it would not be surprising if some of you see that this interval covers zero.

What Did Learn From Q-learning?

At the beginning of school year

IF medication in prior year = {NO}

THEN stage 1 = {BMOD}.

ELSE IF medication in prior year ={YES}

THEN stage 1 = {MED}

Then, every month,

beginning at week 8

• • •

What Did Learn From Q-learning?

… *Then, every month,*
beginning at week 8

IF response status to Stage 1 = {NR}
Then,
IF adherence to MED or BMOD= {NO},
THEN Stage 2 = {AUGMENT}.
Else IF adherence to MED = {YES},
THEN Stage 2 = {INTENSIFY}.
Else IF adherence to BEMOD = {YES},
THEN Stage 2 = {AUGMENT} or {ENHANCE}.
ELSE IF **response status** to Stage 1 = {R}
Then, CONTINUE Stage 1.

What Did Learn From Q-learning?

- The mean Y, school performance, under the more deeply tailored AI obtained via Q-learning is estimated to be **3.72**.
- As expected, this is larger than the value of the AI that started with BMOD and used AUGMENT for non-responders (mean = **3.51**)
- Recall (BMOD, AUGMENT) was the AI with the largest mean among the 4 embedded AIs.

The SAS code you received with this workshop shows you code for how to get the mean under the more deeply-tailored AI discovered by Q-Learning. We do not have time to go over this in this Module.

References

- Nahum-Shani, I., Qian, M., Almirall, D., Pelham, W. E., Gnagy, B., Fabiano, G. A., ... & Murphy, S. A. (2012). Q-learning: A data analysis method for constructing adaptive interventions. Psychological methods, 17(4), 478.

Practicum

- On your return airplane flight, practice Q-Learning in SAS with the simulated Autism study data. As with the previous modules, do this by filling in the ??? in the SAS starter file!

To end the workshop (70min practicum):

- Review and revise your SMART design in your field
- Revise the outline of your specific aims page providing rationale for your SMART

- Individual work (10min)
  Group work (25min)
  Group discussion (35min): 2-3 new groups volunteer

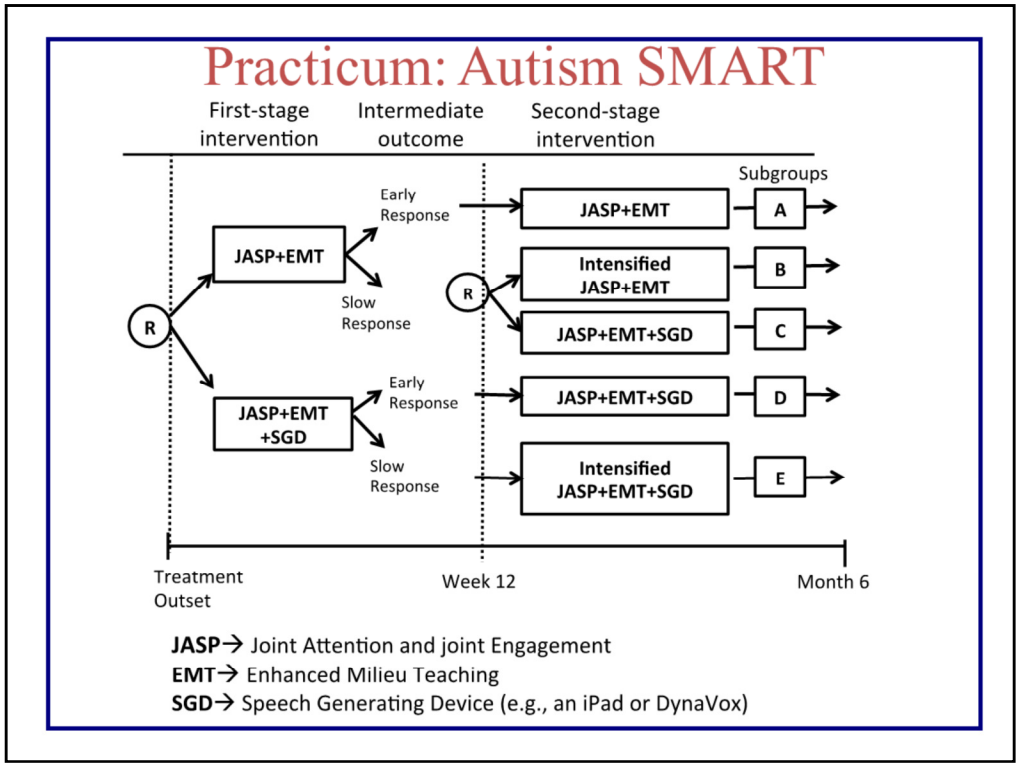

We are now going to practice all of our new data analysis skills using a new data set based on an AUTISM SMART that is still currently in the field. Y

You have a handout with this design printed on it. Keep this handout while we go through the practicum.

Extra Slides: For Statistical Adventure !

- In the next two slides, we actually do Steps 2 + 3 of Q-Learning by hand.
- This is what the PROC QLEARN software automatically does.
- The issue, however, is that the standard errors here are incorrect.
- PROC QLEARN calculates the appropriate standard errors.
- Page 12 on your SAS code doc file

SAS code for Step 2 manually

```
data dat11;
  set dat1;
  * first, everyone gets their observed outcome;
  yhat = y;
  * second, re-assign the outcome for non-responders;
  if R=0 then
        yhat = 3.0039 - 0.2462*o11c - 0.2961*o12c + 0.0391*o13c
            + 0.4868*o14c + 0.0758*a1 - 0.0097*o21c - 0.0980*o22
            + abs(-0.8640*a2 - 0.1934*a1*a2 + 1.1826*o22*a2) ;
run;
proc means data=dat11; var y yhat; run;
```

The MEANS Procedure

Variable	N	Mean	Std Dev
Y	150	2.9533333	1.2814456
yhat	150	3.4107823	0.9385790

This analysis is with simulated data.

This is Step (ii) of Qlearning (done manually): This SAS code manually assigns yhat to prepare for the Step (iii) regression. The coefficients in the definition of yhat are from the Step (i) model results. The absolute value does algebraically what was shown graphically in the previous slide: that is, it assigns the outcome had the child been assigned their optimal treatment at stage 2 (i.e., assigning ADD vs INTENSITY based on adherence to first-stage treatment).

One way to think about the absolute value is that it defines what would be the contribution of stage 2 to the expected outcome given stage 1 and adherence, if we had assigned the best second-stage for the non-responder

As you can see from the PROC MEANS output, yhat has larger mean than y. This is expected if we did it right!!

Note:
We are doing all of this manually here, but in forthcoming slides we describe SAS software (PROC QLEARN) that does all of this automatically!

SAS code for Step 3 manually

```
* Step 3 regression using the new outcome;

proc genmod data=dat11;
  model yhat = o11c o12c o14c o13 a1*o13 a1;
  estimate 'BMOD vs MED given MED prior yr    ' a1*o13 2 a1 2;
  estimate 'BMOD vs MED given NO MED prior yr' a1*o13 0 a1 2;
run;
* medication in the year prior appears to be a tailoring variable  ;
* however, statistical inferences (p-values, confidence intervals) ;
* should not be based on this output.                                ;
```

Contrast Estimate Results

		Naïve 95% Conf Limits		
Label	Estimate	Lower	Upper	P-value
BMOD vs MED given MED prior yr	-0.5012	-0.9423	-0.0601	0.0259
BMOD vs MED given MED prior yr	0.6208	0.3266	0.9150	<0.001

This analysis is with simulated data.

Step (iii) of Qlearning done manually.

Contrast 1: Among children who had medication (and found it acceptable) in the prior year, then starting with MED is better (negative effect).

Contrast 2: Among children who did not have medication (or did and found it unacceptable) in the prior year, then starting with BMOD is better (positive effect).

HOWEVER, WE CANNOT TRUST THE 95% CONFIDENCE INTERVALS PROVIDED BY THE REGRESSION PROCEDURE HERE !!!