

# **Secondary Aims Using Data Arising from a SMART**

Module 6

# General Objectives

- A taste of how data from a SMART can be analyzed to address various scientific questions
  - How to frame scientific questions
  - Experimental cells to be compared
  - Resources you can use for data analysis
  - Less details, more focus on making you feel comfortable with the general approach.

# Outline

Discuss moderators analysis in the context of a SMART

Discuss the idea of “a more deeply-tailored AI”

Q-Learning

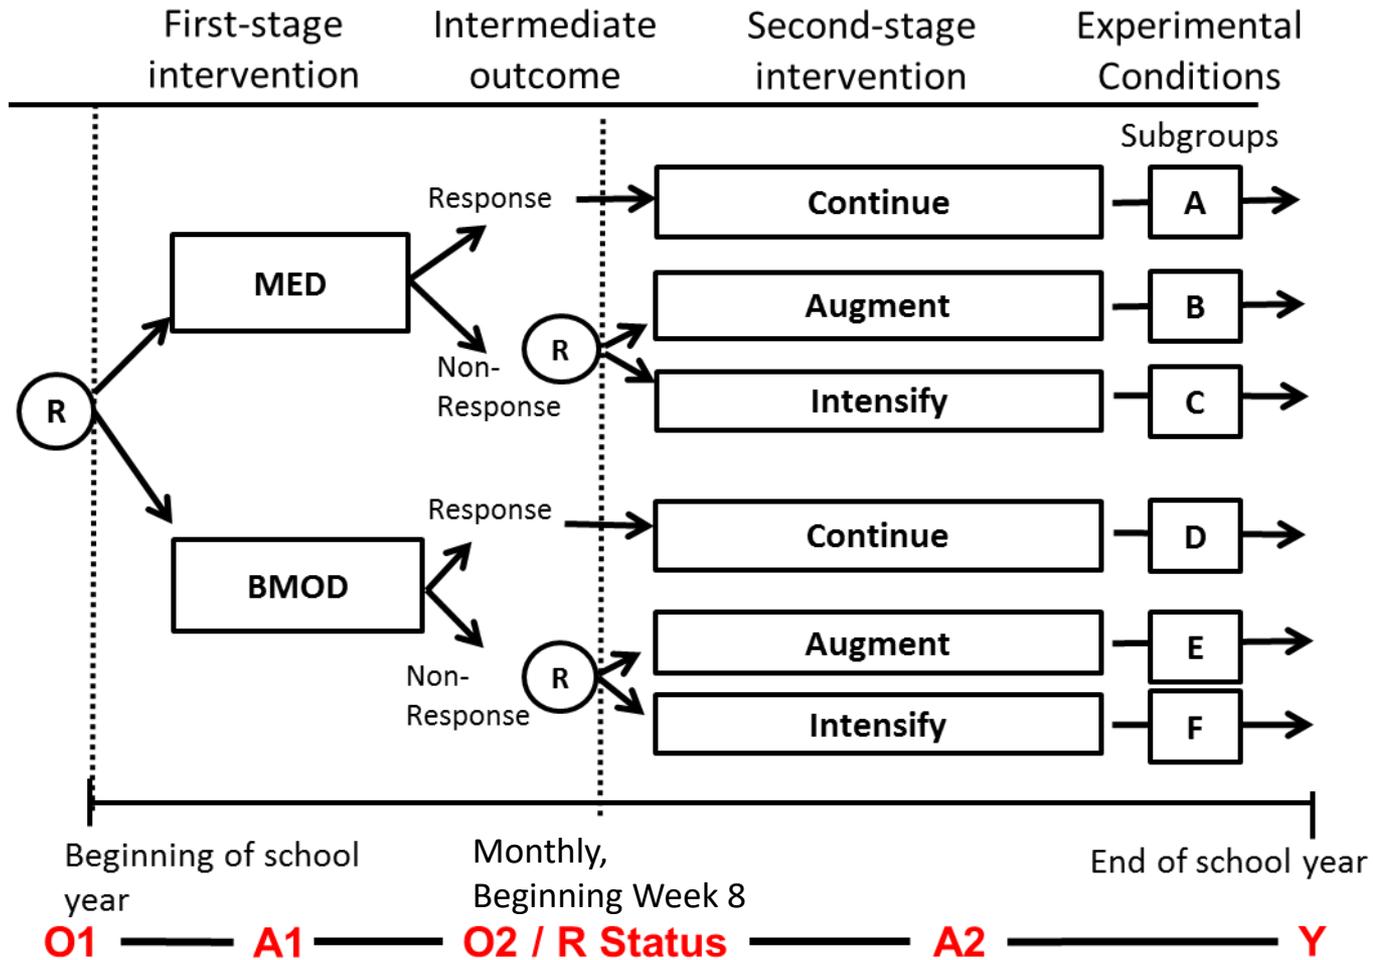
# Outline

## Discuss moderators analysis in the context of a SMART

Discuss the idea of “a more deeply-tailored AI”

Q-Learning

# Remember ADHD SMART?



# ADHD SMART

PI: Pelham

4 embedded adaptive interventions

## ***AI #1:***

Start with MED;  
if non-responder AUGMENT,  
else CONTINUE

## ***AI #2:***

Start with BMOD;  
if non-responder AUGMENT,  
else CONTINUE

## ***AI #3:***

Start with MED;  
if non-responder INTENSIFY,  
else CONTINUE

## ***AI #4:***

Start with BMOD;  
if non-responder INTENSIFY,  
else CONTINUE

# ADHD SMART

PI: Pelham

4 embedded adaptive interventions

## ***AI #1:***

Start with MED;  
if non-responder AUGMENT,  
else CONTINUE

## ***AI #2:***

Start with BMOD;  
if non-responder AUGMENT,  
else CONTINUE

## ***AI #3:***

Start with MED;  
if non-responder INTENSIFY,  
else CONTINUE

## ***AI #4:***

Start with BMOD;  
if non-responder INTENSIFY,  
else CONTINUE

## Moderator analyses dive deeper...

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 also benefit more from AUGMENT vs. INTENSIFY

*For example:* those who do not adhere to initial treatment

*These analyses may suggest new tailoring variables that we should use in our AI.*

# Outline

Discuss moderators analysis in the context of a SMART

**Discuss the idea of “a more deeply-tailored AI”**

Q-Learning

# What is a more deeply-tailored AI?

In the ADHD SMART, there is only 1 tailoring variable embedded *by design* in the embedded AIs

## **AI #1:**

Start with MED;  
if **non-responder** AUGMENT,  
else CONTINUE

## **AI #2:**

Start with BMOD;  
if **non-responder** AUGMENT,  
else CONTINUE

## **AI #3:**

Start with MED;  
if **non-responder** INTENSIFY,  
else CONTINUE

## **AI #4:**

Start with BMOD;  
if **non-responder** INTENSIFY,  
else CONTINUE

# What is a more deeply-tailored AI?

A more **deeply tailored AI** is a sequence of decision rules that include tailoring variables **beyond those** embedded in the SMART *by design*.

- i.e., an AI that tailors treatment to **Response Status** **AND** additional variables

For example, I want to investigate whether to tailor based on:

- First stage MED vs. BMOD on prior receipt of medication; and
- Second stage INT vs. AUG on first-stage adherence.

# The embedded AI looked like this:

*At the beginning of school year,*

**Stage 1 = {BMOD}.**

*Then, every month,*

*beginning at week 8*

**IF response status** to Stage 1 = {NR}

**THEN** Stage 2 = {AUGMENT}.

**ELSE CONTINUE** Stage 1.

# A (hypothetical) 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}

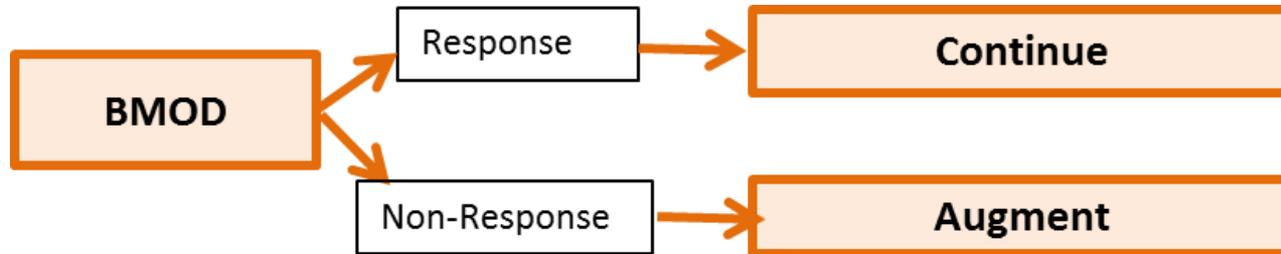
THEN IF **adherence** to stage 1 = {NO},

THEN Stage 2 = {AUGMENT}.

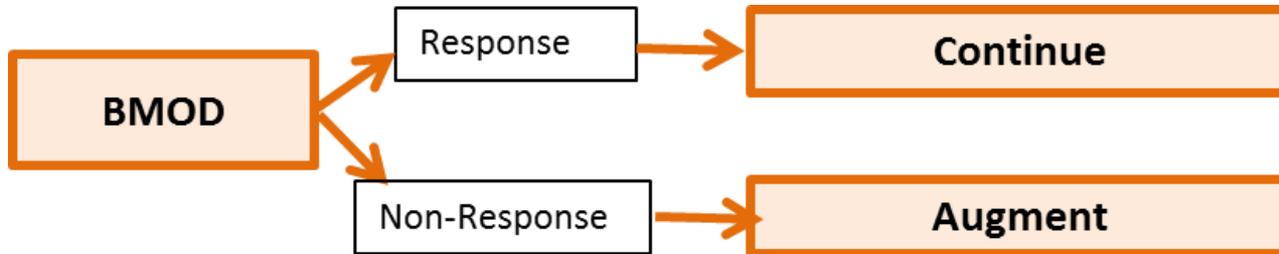
ELSE Stage 2 = {AUGMENT} or {INTENSIFY}.

ELSE CONTINUE Stage 1.

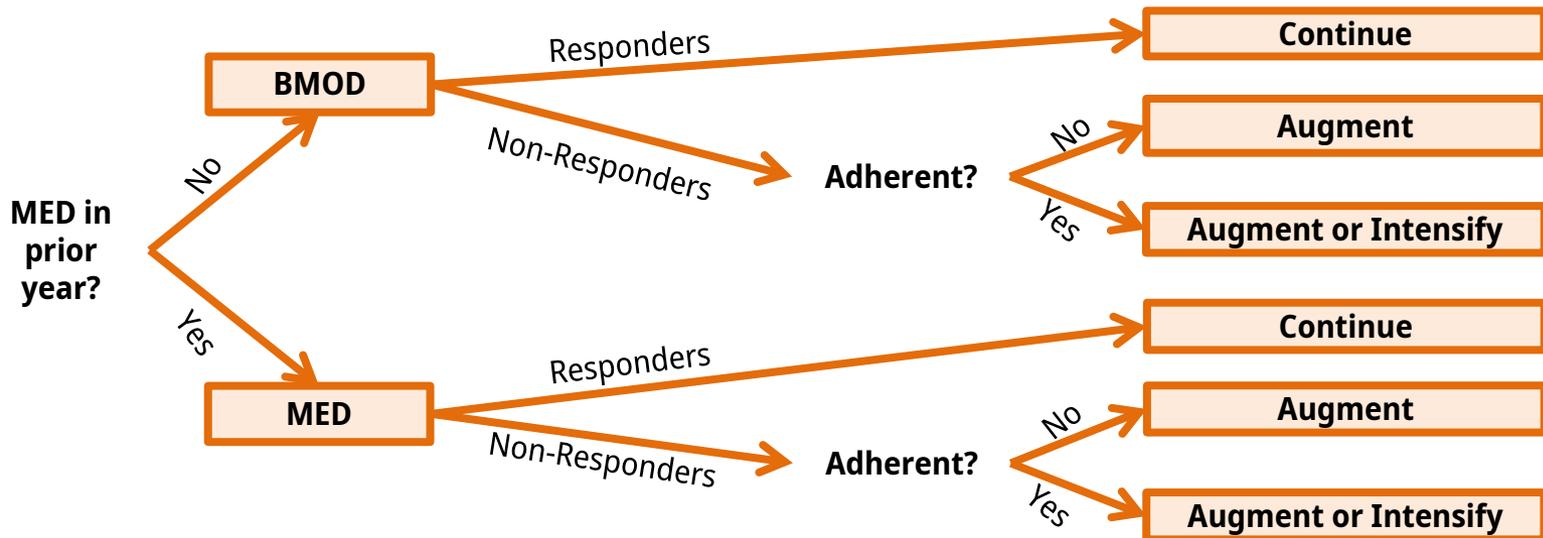
# An Embedded AI:



# An Embedded AI:



# A More Deeply-Tailored AI:



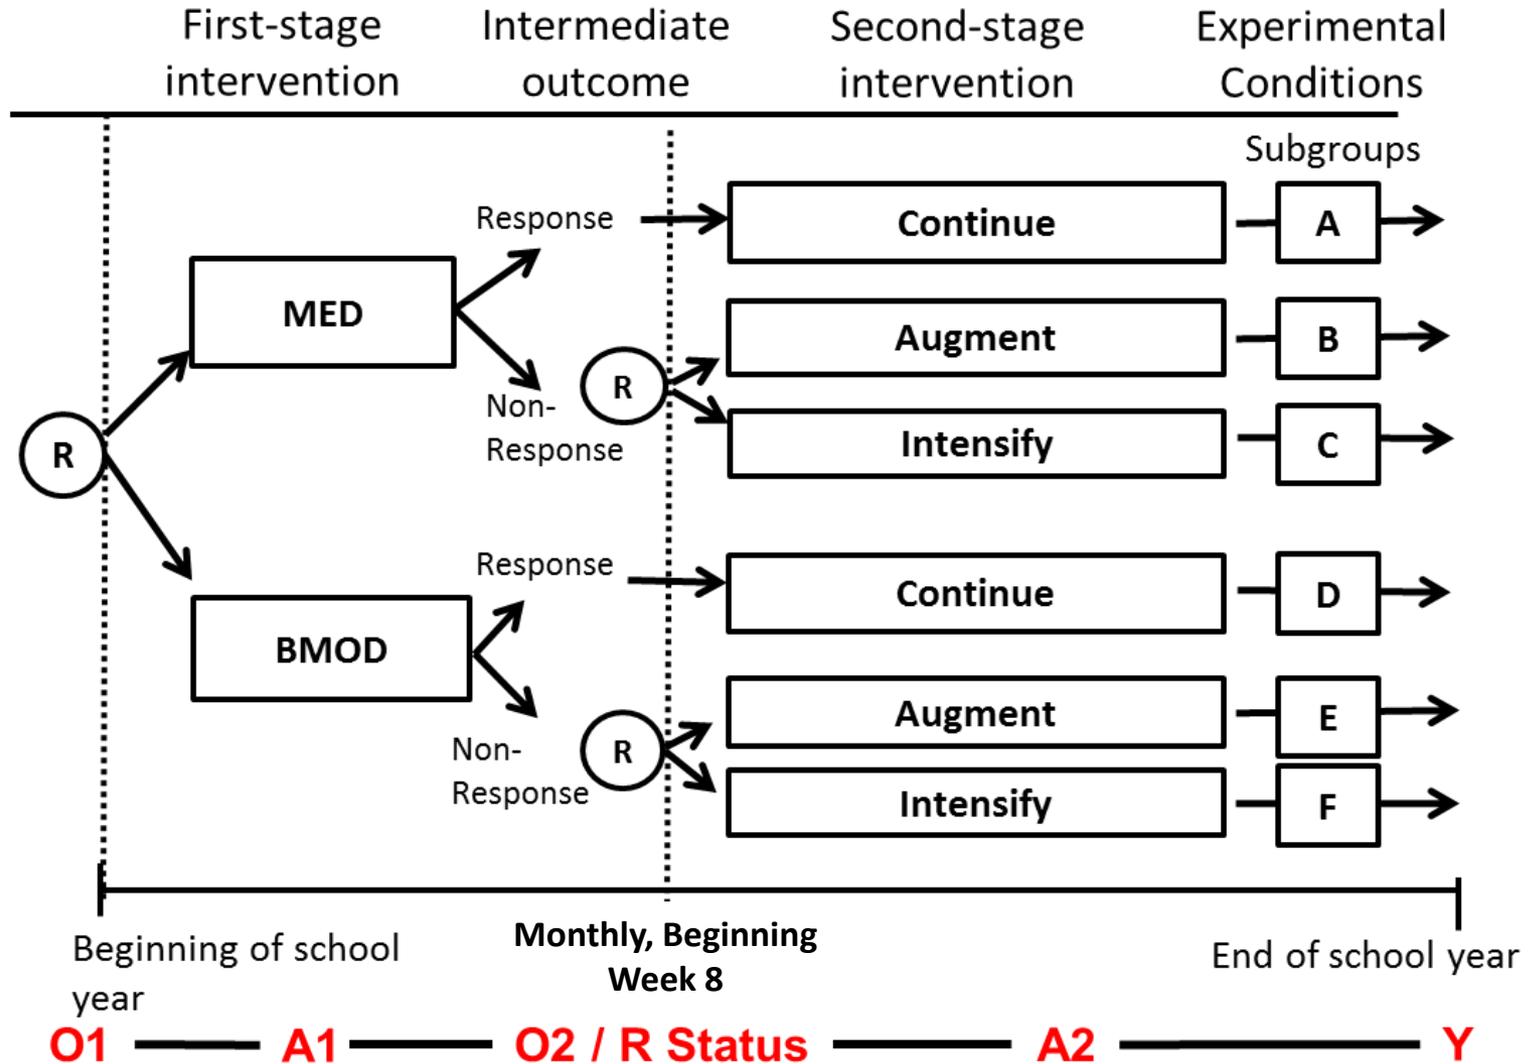
# Outline

Discuss moderators analysis in the context of a SMART

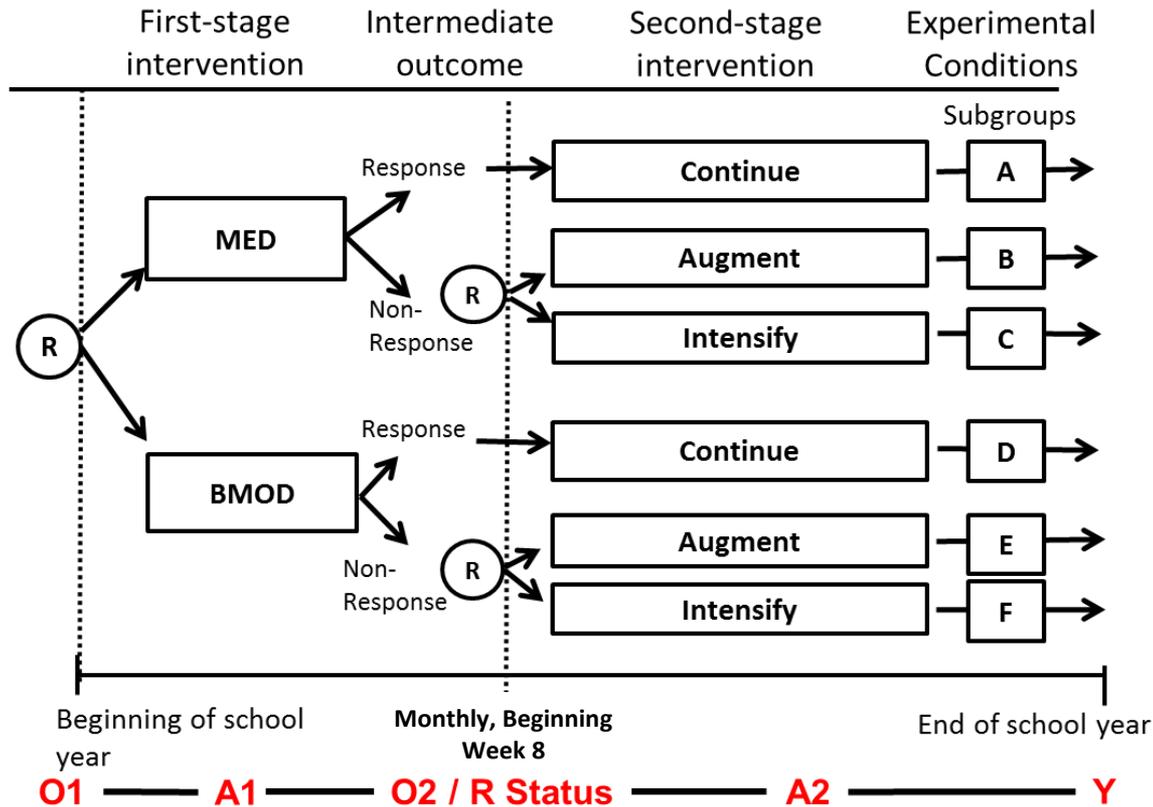
Discuss the idea of “a more deeply-tailored AI”

## Q-Learning

# Other measures collected in a SMART

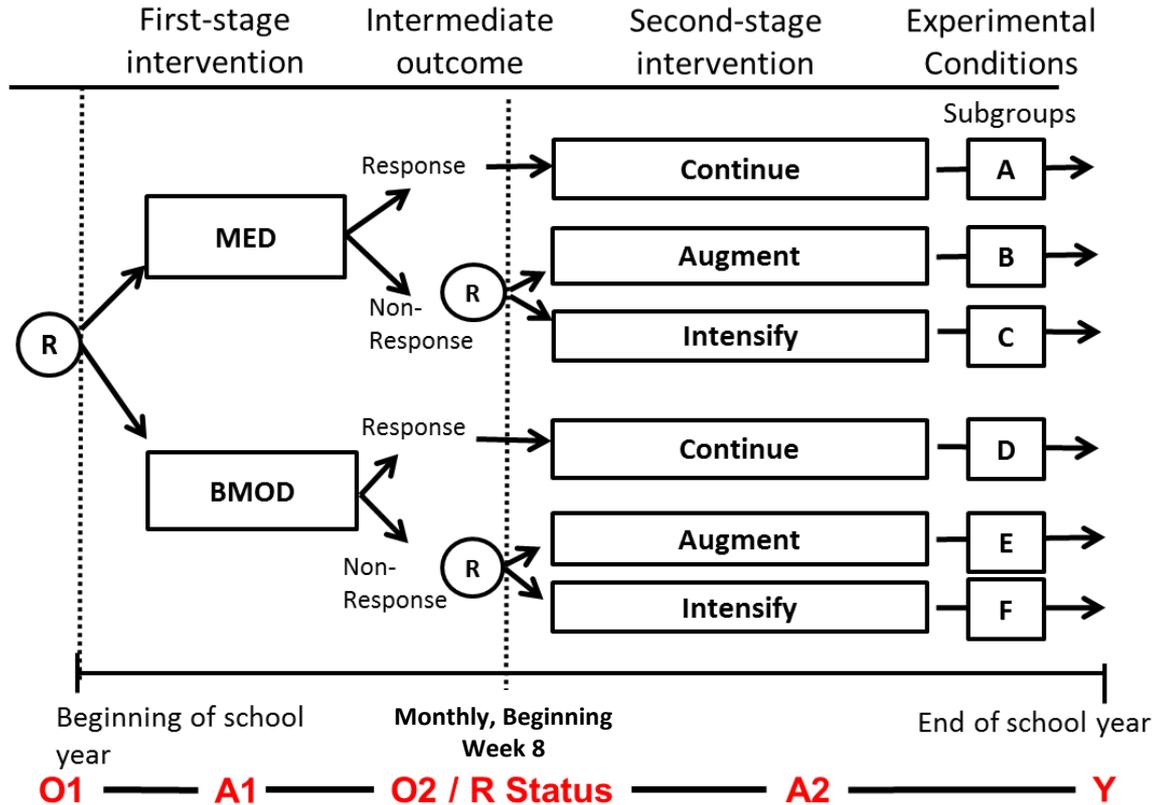


# Other measures collected in a SMART



O1 = Demographics , Med before stage 1, Baseline ADHD scores, Baseline school performance, ODD, ...

# Other measures collected in a SMART



O1 = Demographics , Med before stage 1, Baseline ADHD scores, Baseline school performance, ODD, ...

O2 = Month of non-response, adherence to stage 1, parent functioning during stage 1

## How should we use O1 and O2?

Auxiliary data from **O1** can help decide **who would benefit more from MED vs. BMOD**

- *Our example:* Medication in the prior year

Auxiliary data from **O1 and O2** can help decide **which *non-responders* would benefit more from INTENSIFY vs. AUGMENT**

- *In addition* to using information on first-stage treatment assignment
- *Our example:* Adherence to Stage One treatment

# How do we do baseline moderators analysis?

If the goal is to examine only baseline moderators,  $O_1$ , one could use a single regression:

$$E[Y \mid O_1, A_1, A_2] = \beta_0 + \beta_1 O_1 + \beta_2 A_1 + \beta_3 O_1 A_1 + \beta_4 A_2 + \beta_5 A_1 A_2$$

This would allow us to examine baseline **O1** variables as moderators of first-stage treatment.

# Baseline and time-varying moderators analysis?

If the goal is to examine baseline and time-varying moderators, the instinct might be to do this in a single regression as follows:

$$E[Y \mid \mathbf{O}_1, A_1, \mathbf{O}_2, A_2] = \beta_0 + \beta_1 \mathbf{O}_1 + \beta_2 A_1 + \beta_3 \mathbf{O}_1 A_1 + \beta_4 \mathbf{O}_2 + \beta_5 A_2 + \beta_6 \mathbf{O}_2 A_2 + \beta_7 A_1 A_2$$

The hope is that this would allow us to examine baseline **O1** variables as moderators of first-stage treatment & time-varying **O2** variables as moderators of second-stage treatment.

*What's the problem with this approach?*

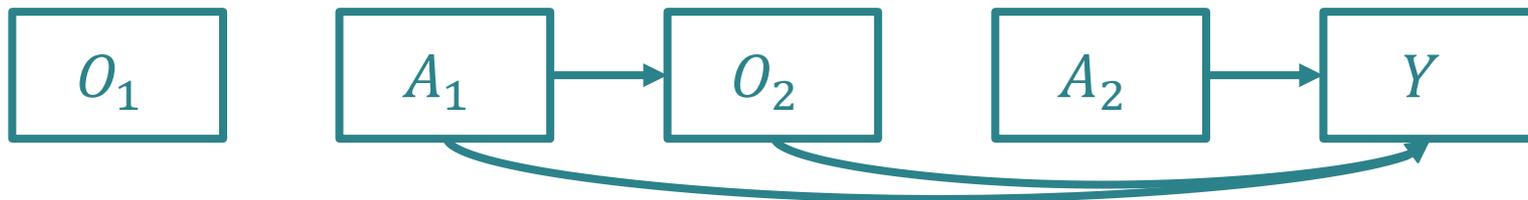
# Baseline and time-varying moderators analysis?

Instinct might be to do this in a single regression:

$$E[Y | \mathbf{O}_1, A_1, \mathbf{O}_2, A_2] = \beta_0 + \beta_1 \mathbf{O}_1 + \beta_2 A_1 + \beta_3 \mathbf{O}_1 A_1 \\ + \beta_4 \mathbf{O}_2 + \beta_5 A_2 + \beta_6 \mathbf{O}_2 A_2 \\ + \beta_7 A_1 A_2$$

**O2** happens after **A1**

- Potential mediator of relationship between A1 and Y



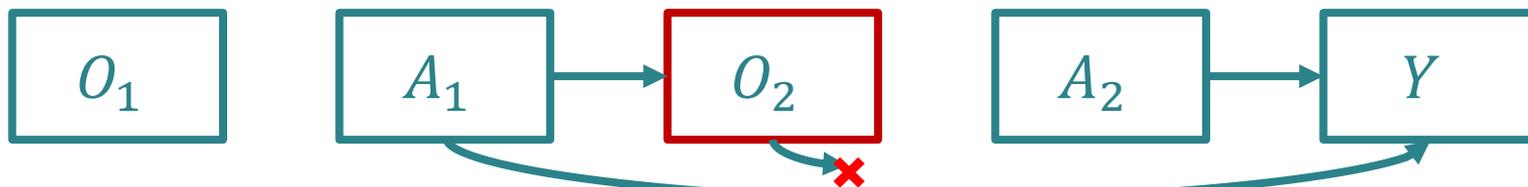
# Baseline and time-varying moderators analysis?

Instinct might be to do this in a single regression:

$$E[Y | \mathbf{O}_1, A_1, \mathbf{O}_2, A_2] = \beta_0 + \beta_1 \mathbf{O}_1 + \beta_2 A_1 + \beta_3 \mathbf{O}_1 A_1 \\ + \beta_4 \mathbf{O}_2 + \beta_5 A_2 + \beta_6 \mathbf{O}_2 A_2 \\ + \beta_7 A_1 A_2$$

**O2** happens *after* **A1**

- Potential mediator of relationship between A1 and Y
- Thus,  $(\beta_2, \beta_3)$  do not have the causal interpretation one wants
  - Because conditioning naively on O2 “cuts-off” the indirect effect of A1 on Y via O2



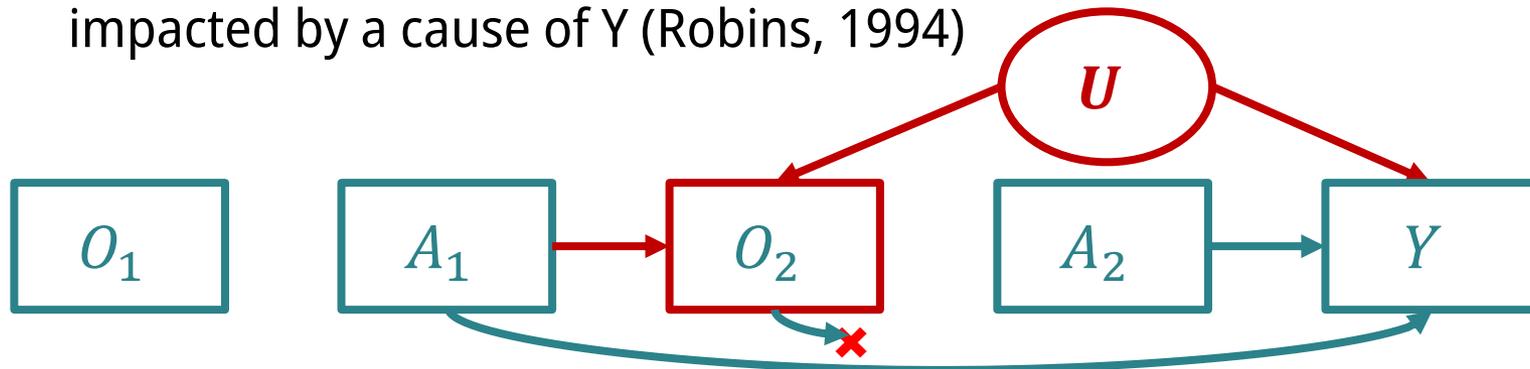
# Baseline and time-varying moderators analysis?

Instinct might be to do this in a single regression:

$$E[Y | \mathbf{O}_1, A_1, \mathbf{O}_2, A_2] = \beta_0 + \beta_1 \mathbf{O}_1 + \beta_2 A_1 + \beta_3 \mathbf{O}_1 A_1 \\ + \beta_4 \mathbf{O}_2 + \beta_5 A_2 + \beta_6 \mathbf{O}_2 A_2 \\ + \beta_7 A_1 A_2$$

**O2** happens *after* **A1**

- Collider Bias: A spurious (non-causal) correlation between A1 and Y resulting from conditioning naively on an outcome of A1 that is also impacted by a cause of Y (Robins, 1994)



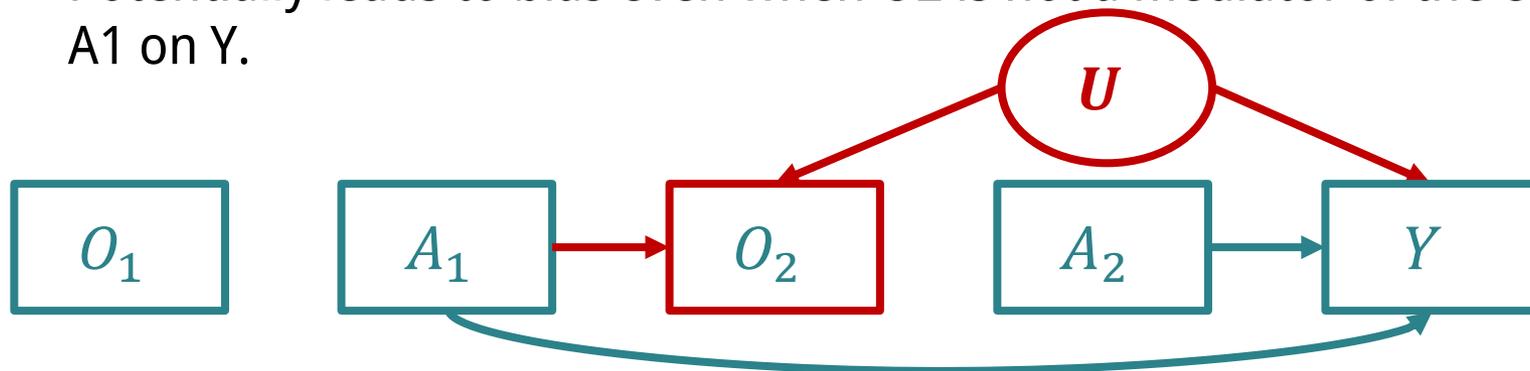
# Baseline and time-varying moderators analysis?

Instinct might be to do this in a single regression:

$$E[Y | \mathbf{O}_1, A_1, \mathbf{O}_2, A_2] = \beta_0 + \beta_1 \mathbf{O}_1 + \beta_2 A_1 + \beta_3 \mathbf{O}_1 A_1 + \beta_4 \mathbf{O}_2 + \beta_5 A_2 + \beta_6 \mathbf{O}_2 A_2 + \beta_7 A_1 A_2$$

**O2** happens after **A1**

- Collider Bias (Robins, 1994)
- Potentially leads to bias even when O2 is not a mediator of the effect of A1 on Y.



# What is Q-Learning?

- Extends regression to **sequential treatments**
- “Q” = “Quality”
- Hypothesis-generating

# What is Q-Learning?

The results of Q-Learning *propose an AI* with greater treatment individualization

- i.e., an AI that includes more tailoring variables than the AIs embedded in the SMART *by design*

# Q-Learning has 3 steps

## Step 1

- ***Second-Stage Regression***
  - Are **O1, A1, and O2** useful in making decisions about second-stage tactics?
  - (Are **O1, A1, and O2** useful in deciding **which NR** would benefit from Augment vs. Intensify?)

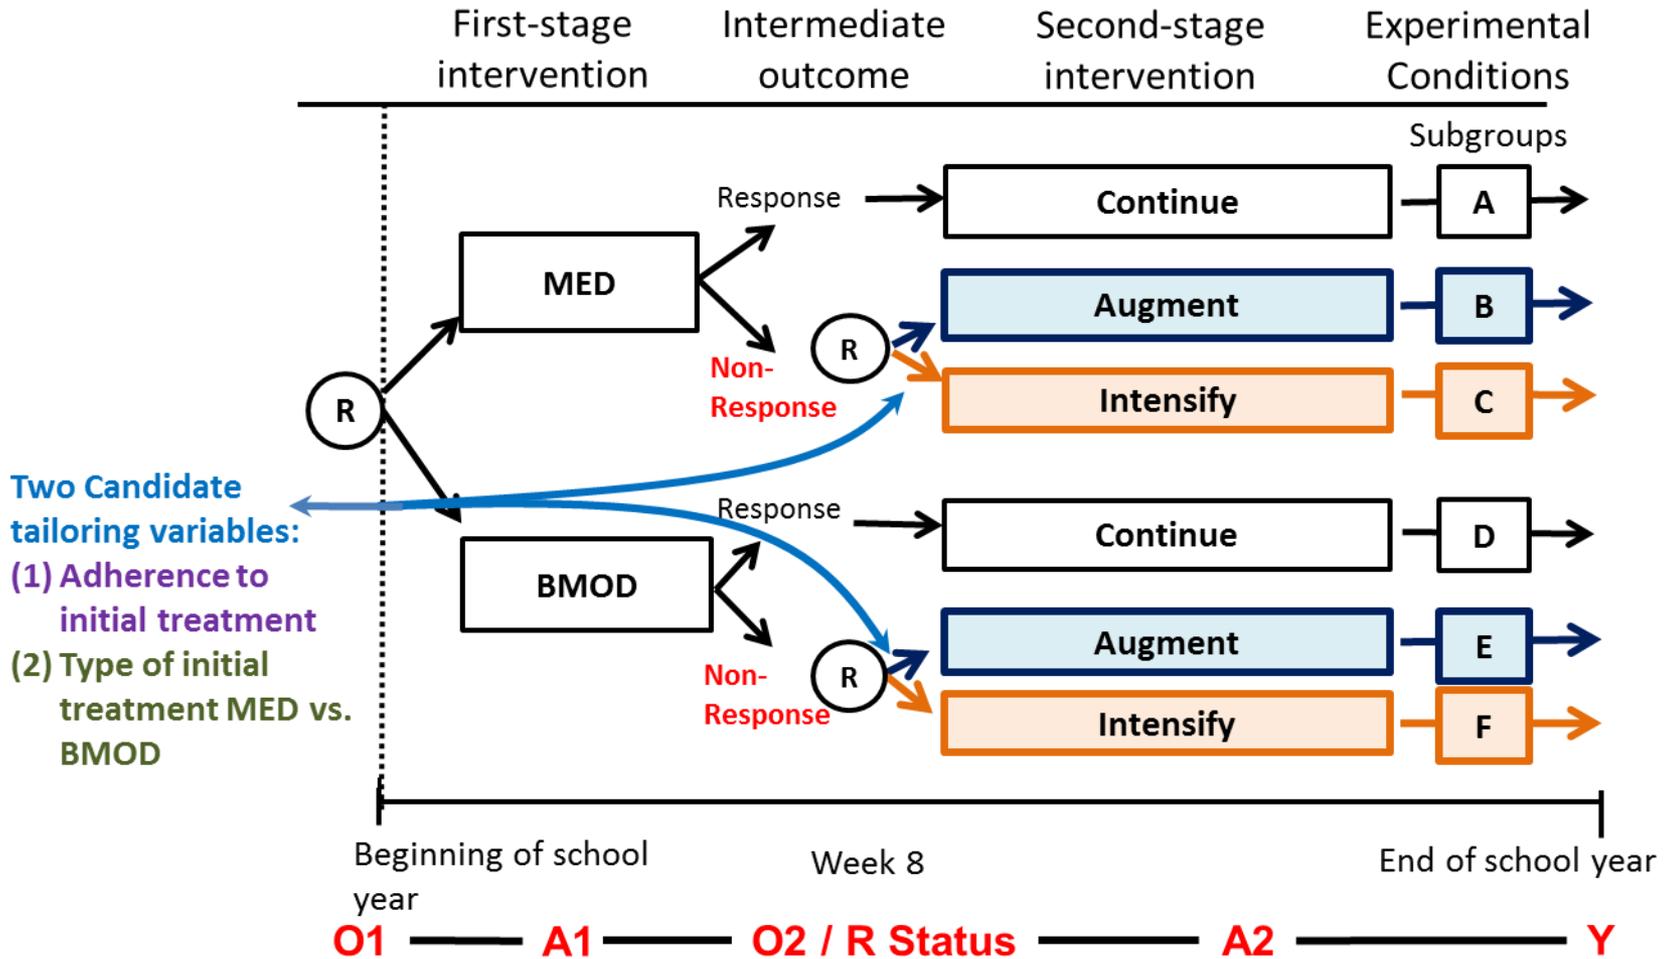
## Step 2

- ***Calculate  $\hat{Y}_i$*** 
  - What would the outcome be if they had received the best second-stage tactic given **O1, A1, and O2**?
  - **$\hat{Y}_i$  is the estimated optimal outcome under the best second-stage tactic for non-responders.** ( $\hat{Y}_i=Y$  for responders)

## Step 3

- ***First-Stage Regression***
  - Is **O1** useful in making decisions about first-stage tactics, *assuming we use optimal second-stage tactic*? (Use  $\hat{Y}_i$  from Step 2 as the outcome!)
  - (Is **O1** useful in deciding who would benefit from MED vs. BMOD, *assuming NRs get the best second-stage treatment*?)

# Step 1: Second-stage tailoring



# Step 1: Second-stage tailoring

In this step, we want 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*?

# Step 1: Second-stage tailoring

In this step, we want 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*?

*To do this:* Fit a **moderators analysis** using data from non-responders

See next slide for details...

# Step 1: Second-stage tailoring

In this step, we want 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*?

To do this: Fit a **moderated regression model** using data from non-responders

$$\begin{aligned} E[Y | O_1, A_1, O_2, A_2] \\ = \beta_0 + \beta_1 O_{11c} + \beta_2 O_{12c} + \beta_3 O_{13c} + \beta_4 O_{14c} + \beta_5 O_{12c} \\ + \beta_6 O_{21c} + \beta_7 \mathbf{A}_1 + \beta_8 \mathbf{O}_{22} \\ + \beta_9 A_2 + \beta_{10} (A_2 \times \mathbf{A}_1) + \beta_{11} (A_2 \times \mathbf{O}_{22}) \end{aligned}$$

## Step 1: Second-stage tailoring

$$E[Y | O_1, A_1, O_2, A_2] = \beta_0 + \dots + \beta_7 \mathbf{A}_1 + \beta_8 \text{adherence} \\ + \beta_9 A_2 + \beta_{10} (A_2 \times \mathbf{A}_1) + \beta_{11} (A_2 \times \text{adherence})$$

**A1 = Stage 1 options: -1=MED; 1=BMOD**

**A2 = Stage 2 options: -1=ADD; 1=INTSFY**

**Adherence to Stage 1: 1=yes; 0=no**

Y = End of year school performance

## Step 1: Second-stage tailoring

$$E[Y | O_1, A_1, O_2, A_2] = \beta_0 + \dots + \beta_7 A_1 + \beta_8 \text{adherence} + \beta_9 A_2 + \beta_{10} (A_2 \times A_1) + \beta_{11} (A_2 \times \text{adherence})$$

**A1 = Stage 1 options: -1=MED; 1=BMOD**

**A2 = Stage 2 options: -1=ADD; 1=INTSFY**

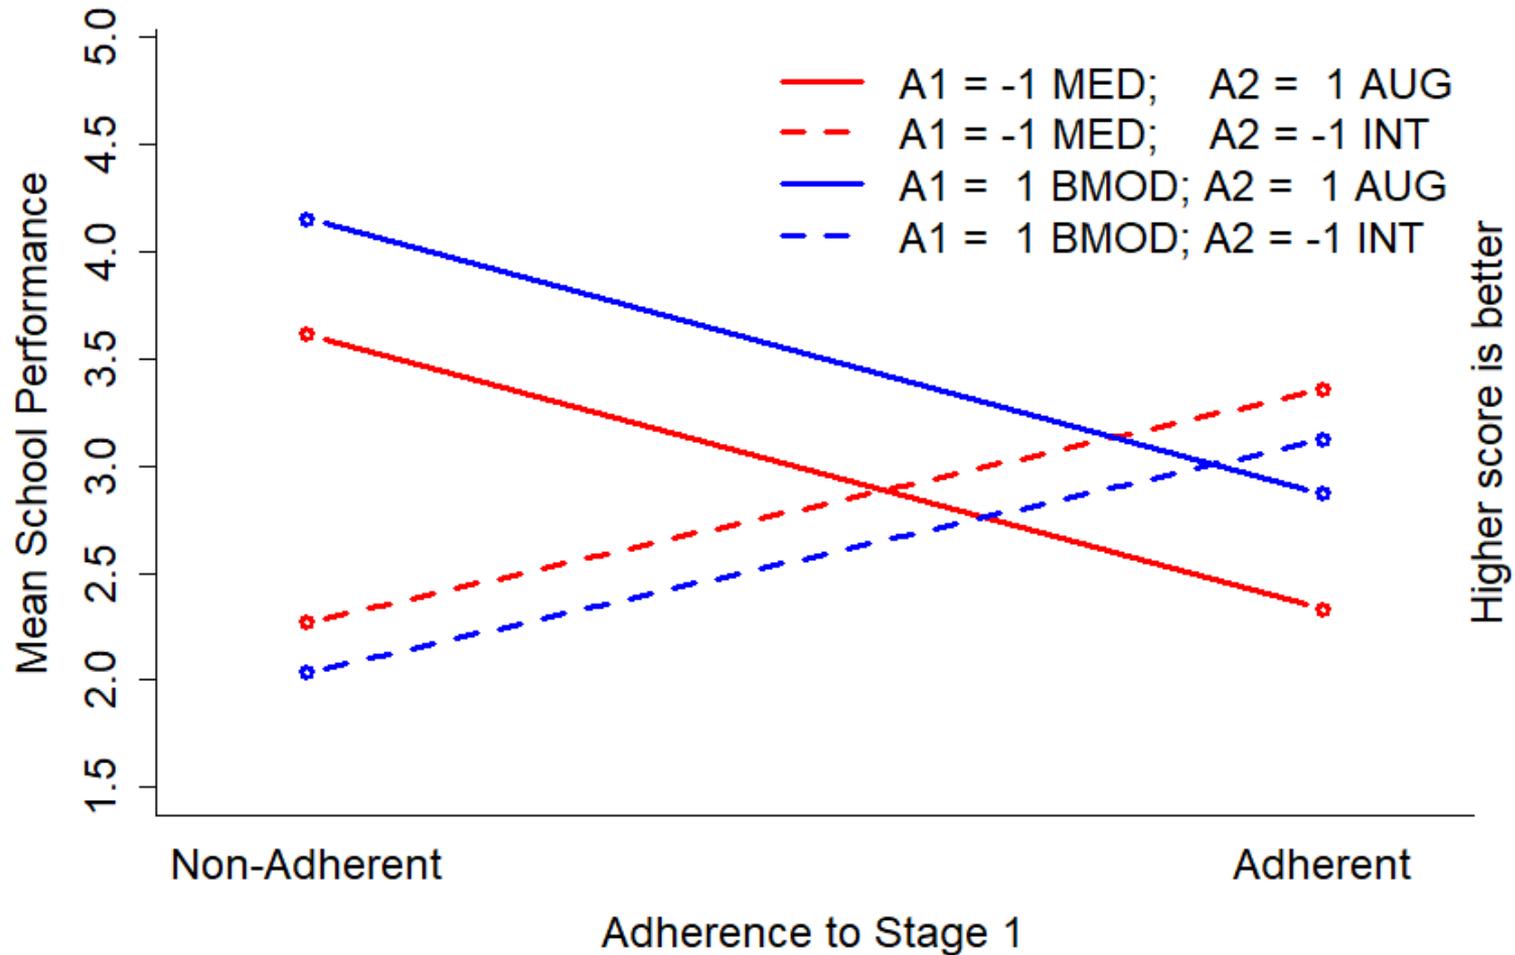
**Adherence to Stage 1: 1=yes; 0=no**

Y = End of year school performance

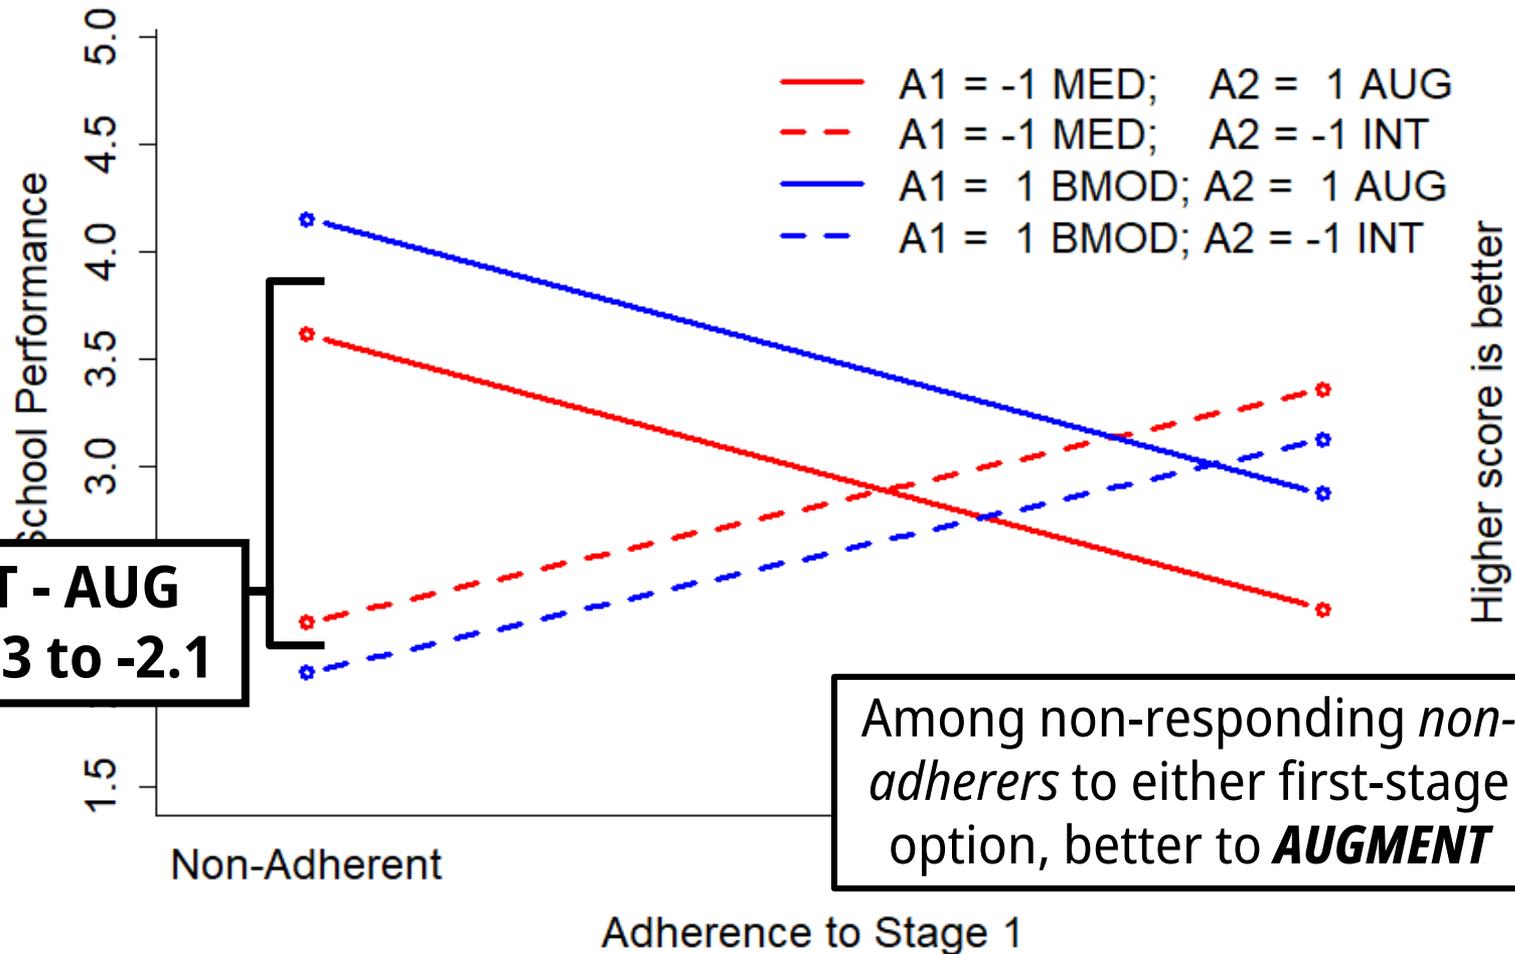
***This model will help us to:***

- Determine whether the best second stage tactics varies depending on the tailoring variables; and
- Identify the best second-stage tactic for each level of the tailoring variable

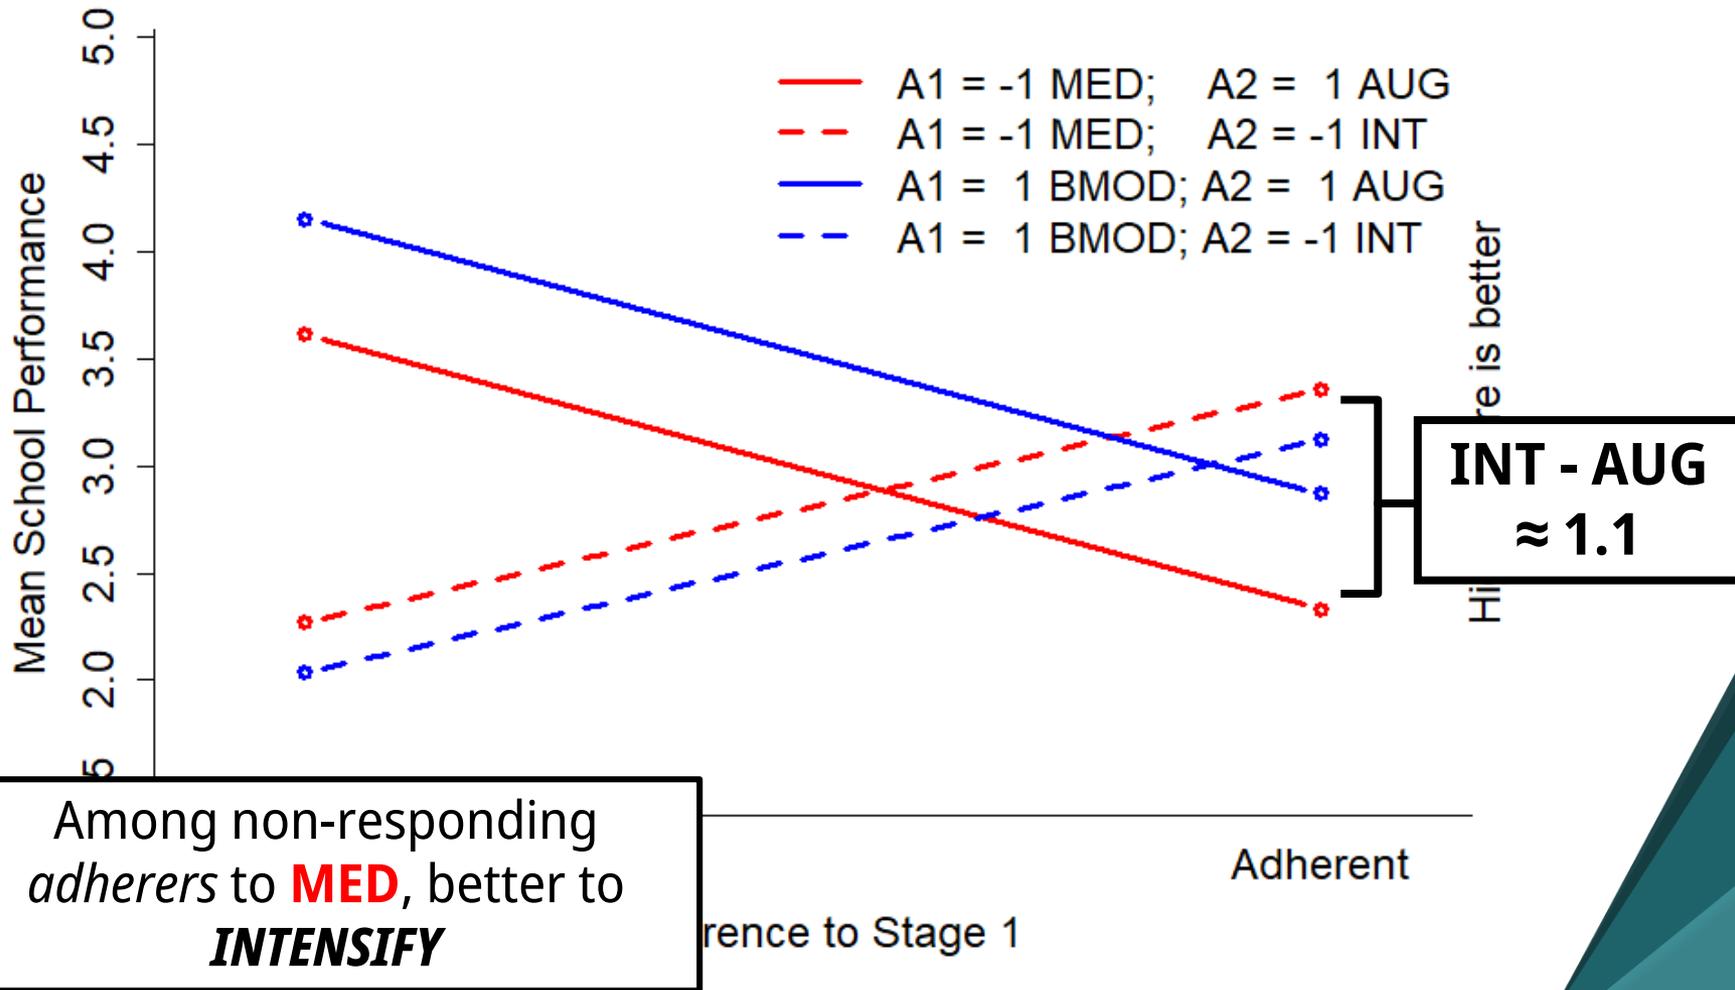
# Step 1: Second-stage tailoring



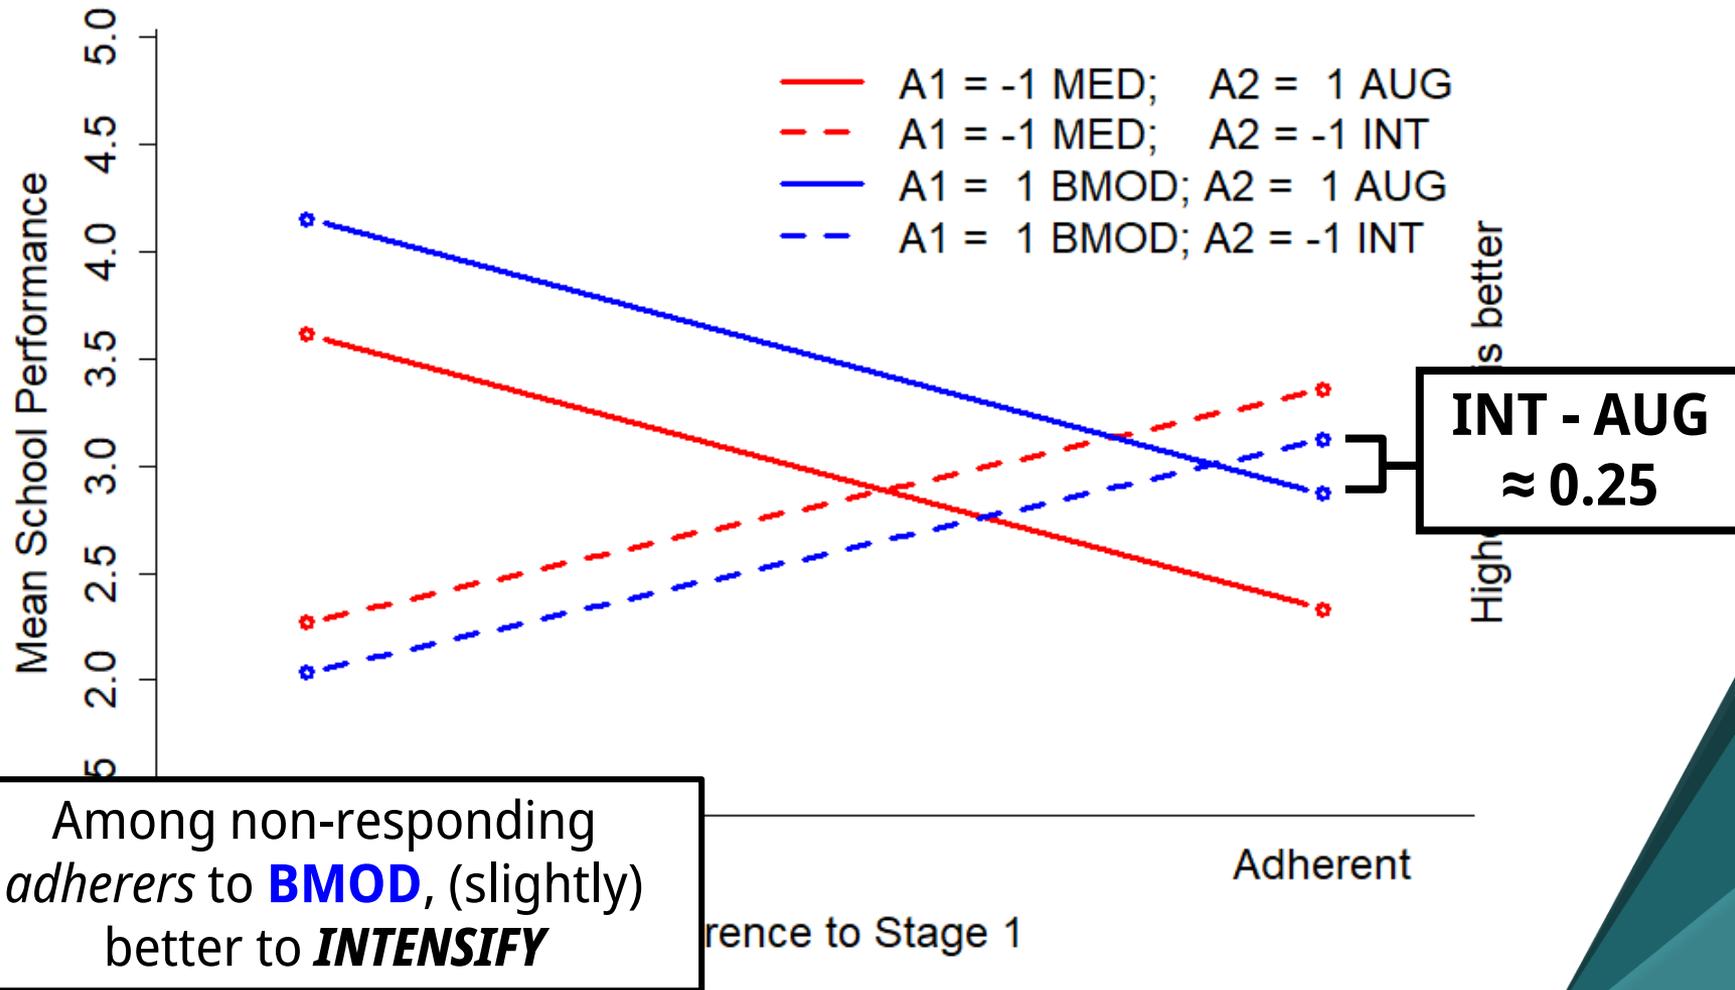
# Step 1: Second-stage tailoring



# Step 1: Second-stage tailoring



# Step 1: Second-stage tailoring



## Step 2: Predicted outcome under the best stage 2 option

Next, we assign each non-responder the value  $\hat{Y}_i$

$\hat{Y}_i$  = The expected outcome if each non-responder received **the best second-stage tactic** given their **initial treatment** and **adherence**

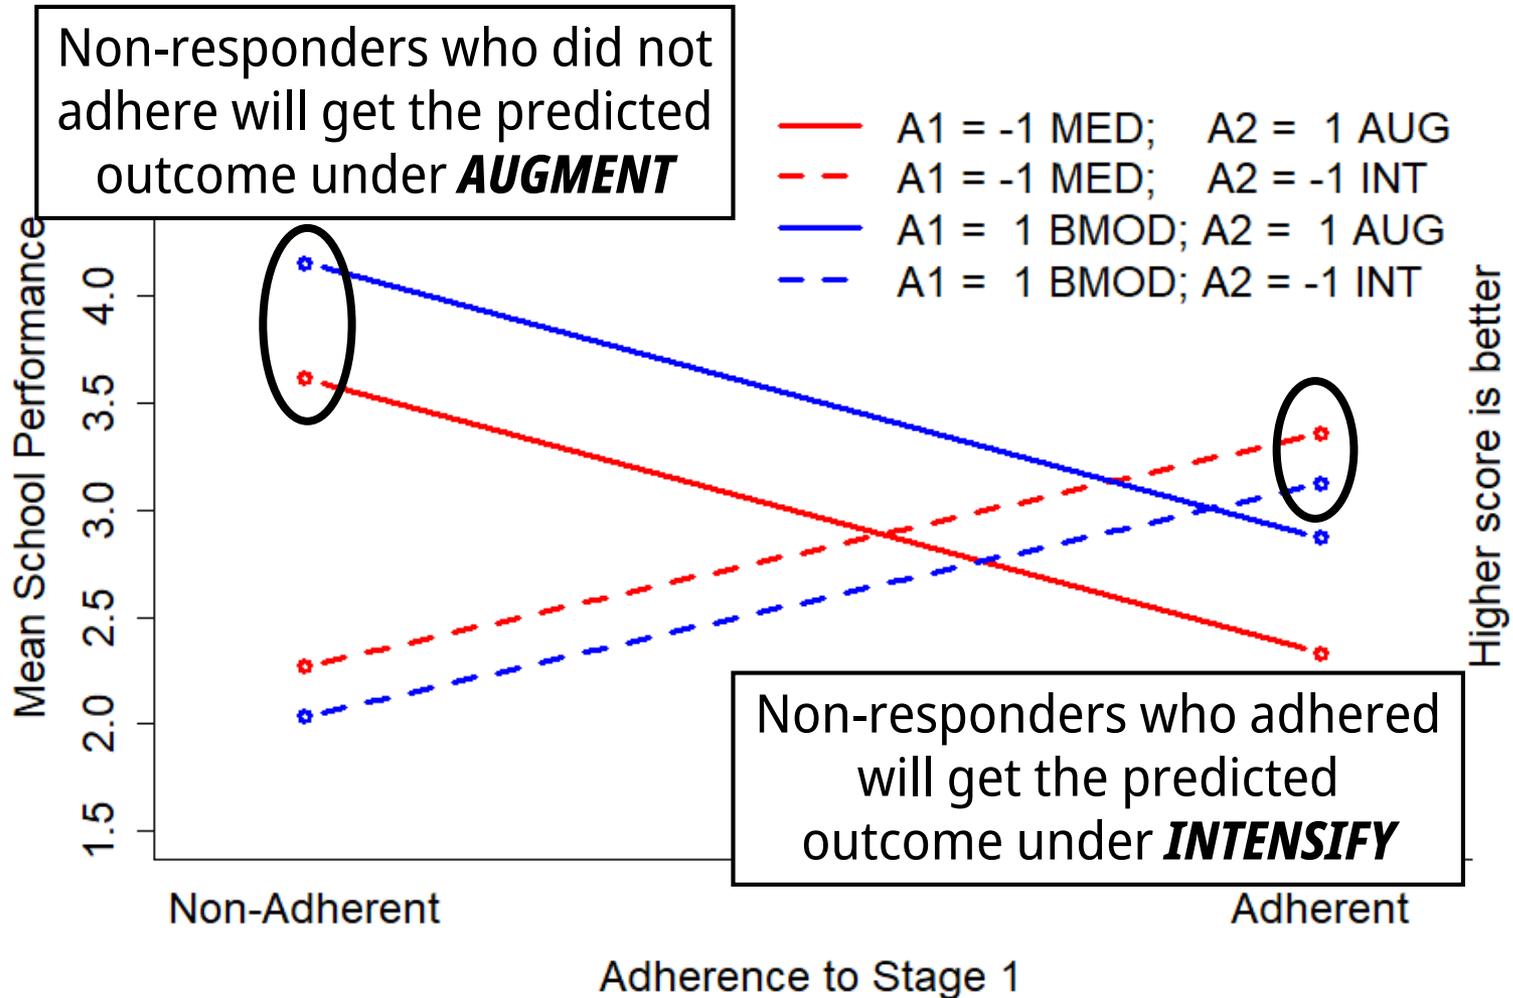
## Step 2: Predicted outcome under the best stage 2 option

Next, we assign each non-responder the value  $\hat{Y}_i$

$\hat{Y}_i$  = The expected outcome if each non-responder received **the best second-stage tactic** given their **initial treatment** and **adherence**

- We used Step 1 Regression to identify the best Stage 2 tactic for any given level of the tailoring variables.
- We use these results to estimate what the outcome for each non-responder if they received the best Stage 2 tactic, *given their observed values* on the tailoring variables
  - Given his/her observed values on the tailoring variables
- Responders?  $\hat{Y}_i = Y_i$

## Step 2: Predicted outcome under the best stage 2 option



## Step 2: Predicted outcome under the best stage 2 option

$$E[Y \mid O_1, A_1, O_2, A_2] = \beta_0 + \dots + \beta_7 \mathbf{A}_1 + \beta_8 \text{adherence} \\ + \beta_9 A_2 + \beta_{10} (A_2 \times \mathbf{A}_1) + \beta_{11} (A_2 \times \text{adherence})$$

**PRETEND FOR A MOMENT THAT:**

$$\beta_0 = 3, \quad \beta_6 = 0.1, \quad \beta_7 = -0.1, \quad \beta_8 = -1, \quad \beta_9 = -0.2, \quad \beta_{10} = 1.2$$

## Step 2: Predicted outcome under the best stage 2 option

$$E[Y | O_1, A_1, O_2, A_2] = \beta_0 + \dots + \beta_7 \mathbf{A}_1 + \beta_8 \text{adherence} \\ + \beta_9 A_2 + \beta_{10} (A_2 \times \mathbf{A}_1) + \beta_{11} (A_2 \times \text{adherence})$$

### PRETEND FOR A MOMENT THAT:

$$\beta_0 = 3, \quad \beta_6 = 0.1, \quad \beta_7 = -0.1, \quad \beta_8 = -1, \quad \beta_9 = -0.2, \quad \beta_{10} = 1.2$$

**John** was a non-responding, non-adhering (adherence=0) participant who received:  
Stage 1: MED (A1 = 1)  
Stage 2: INT (A2 = -1)  
and had mean values for all baseline variables

## Step 2: Predicted outcome under the best stage 2 option

$$E[Y | O_1, A_1, O_2, A_2] = \beta_0 + \dots + \beta_7 \mathbf{A}_1 + \beta_8 \text{adherence} \\ + \beta_9 A_2 + \beta_{10} (A_2 \times \mathbf{A}_1) + \beta_{11} (A_2 \times \text{adherence})$$

**PRETEND FOR A MOMENT THAT:**

$$\beta_0 = 3, \quad \beta_6 = 0.1, \quad \beta_7 = -0.1, \quad \beta_8 = -1, \quad \beta_9 = -0.2, \quad \beta_{10} = 1.2$$

**John** was a non-responding, non-adhering (adherence=0) participant who received:  
Stage 1: MED ( $A_1 = 1$ )  
Stage 2: INT ( $A_2 = -1$ )  
and had mean values for all baseline variables

$$\hat{Y} = 3 + 0.1(A_1) - 0.1(\text{adherence}) - 0.1(A_2) - 0.2(A_2 \times A_1) + 1.2(A_2 \times 0)$$

$$\hat{Y}_{A_2=INT} = 3 + 0.1(1) - 0.1(0) - 0.1(1) - 0.2(1 \times 1) + 1.2(1 \times 0) = 2.8$$

John's score under INT (which he received)

$$\hat{Y}_{A_2=AUG} = 3 + 0.1(1) - 0.1(0) - 0.1(-1) - 0.2(-1 \times 1) + 1.2(-1 \times 0) = 3.4$$

John's score under AUG (which he did not receive)

## Step 2: Predicted outcome under the best stage 2 option

$$E[Y | O_1, A_1, O_2, A_2] = \beta_0 + \dots + \beta_7 \mathbf{A}_1 + \beta_8 \text{adherence} \\ + \beta_9 A_2 + \beta_{10} (A_2 \times \mathbf{A}_1) + \beta_{11} (A_2 \times \text{adherence})$$

**PRETEND FOR A MOMENT THAT:**

$$\beta_0 = 3, \quad \beta_6 = 0.1, \quad \beta_7 = -0.1, \quad \beta_8 = -1, \quad \beta_9 = -0.2, \quad \beta_{10} = 1.2$$

**John** was a non-responding, non-adhering (adherence=0) participant who received:  
Stage 1: MED ( $A_1 = 1$ )  
Stage 2: INT ( $A_2 = -1$ )  
and had mean values for all baseline variables

$$\hat{Y} = 3 + 0.1(A_1) - 0.1(\text{adherence}) - 0.1(A_2) - 0.2(A_2 \times A_1) + 1.2(A_2 \times 0)$$

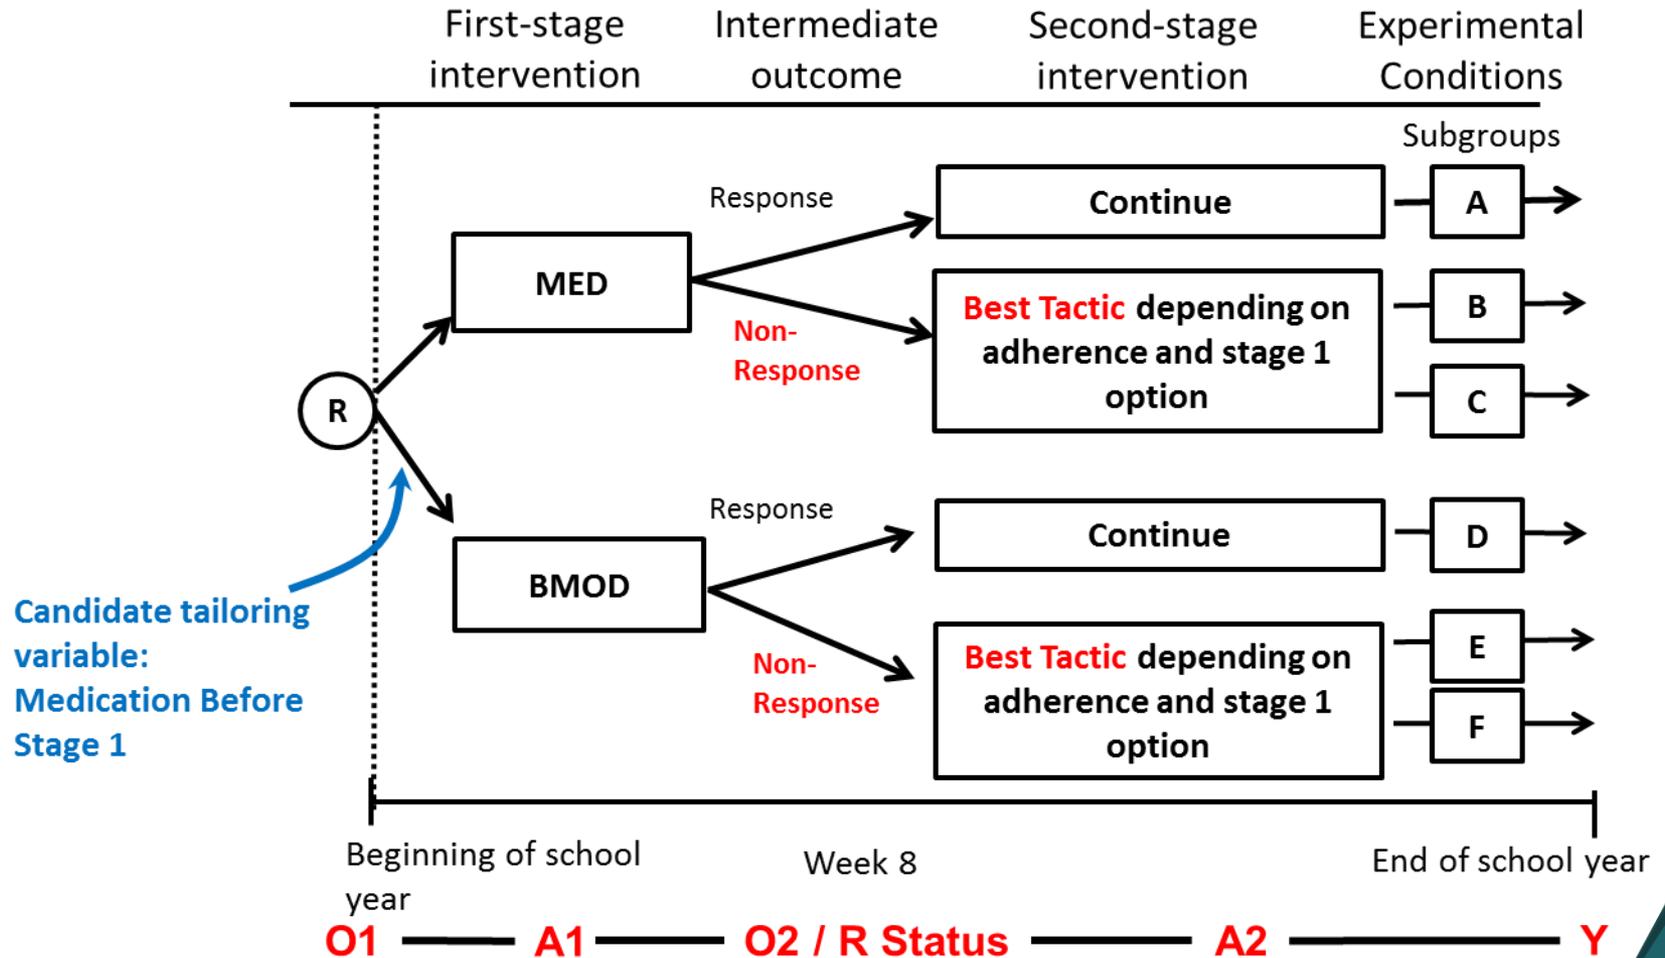
$$\hat{Y}_{A_2=INT} = 3 + 0.1(1) - 0.1(0) - 0.1(1) - 0.2(1 \times 1) + 1.2(1 \times 0) = 2.8$$

John's score under INT (which he received)

$$\hat{Y}_{A_2=AUG} = 3 + 0.1(1) - 0.1(0) - 0.1(-1) - 0.2(-1 \times 1) + 1.2(-1 \times 0) = 3.4$$

**John's  $\hat{Y}_i$**

# 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 to select the *best first-stage option*?
  - Assuming that in the future, non-responders get the best Stage 2 tactic

## 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 to select the best first-stage option?
  - Assuming that in the future, non-responders get the best Stage 2 tactic
- We do this by using  $\hat{Y}_i$  as the outcome in a regression where we explore the usefulness of prior medication for making decisions about first-stage options.

## Step 3: Move backwards to first-stage tailoring

Fit the following regression model:

$$E[\hat{Y} \mid O_1, A_1] = \beta_0 + \beta_1 O_{11c} + \beta_2 O_{12c} + \beta_3 O_{14c} \\ + \beta_4 \text{priorMed} + \beta_5 A_1 + \beta_6 (\text{priorMed} \times A_1)$$

Controlling for  
stage 2 tactic

## Step 3: Move backwards to first-stage tailoring

Fit the following regression model:

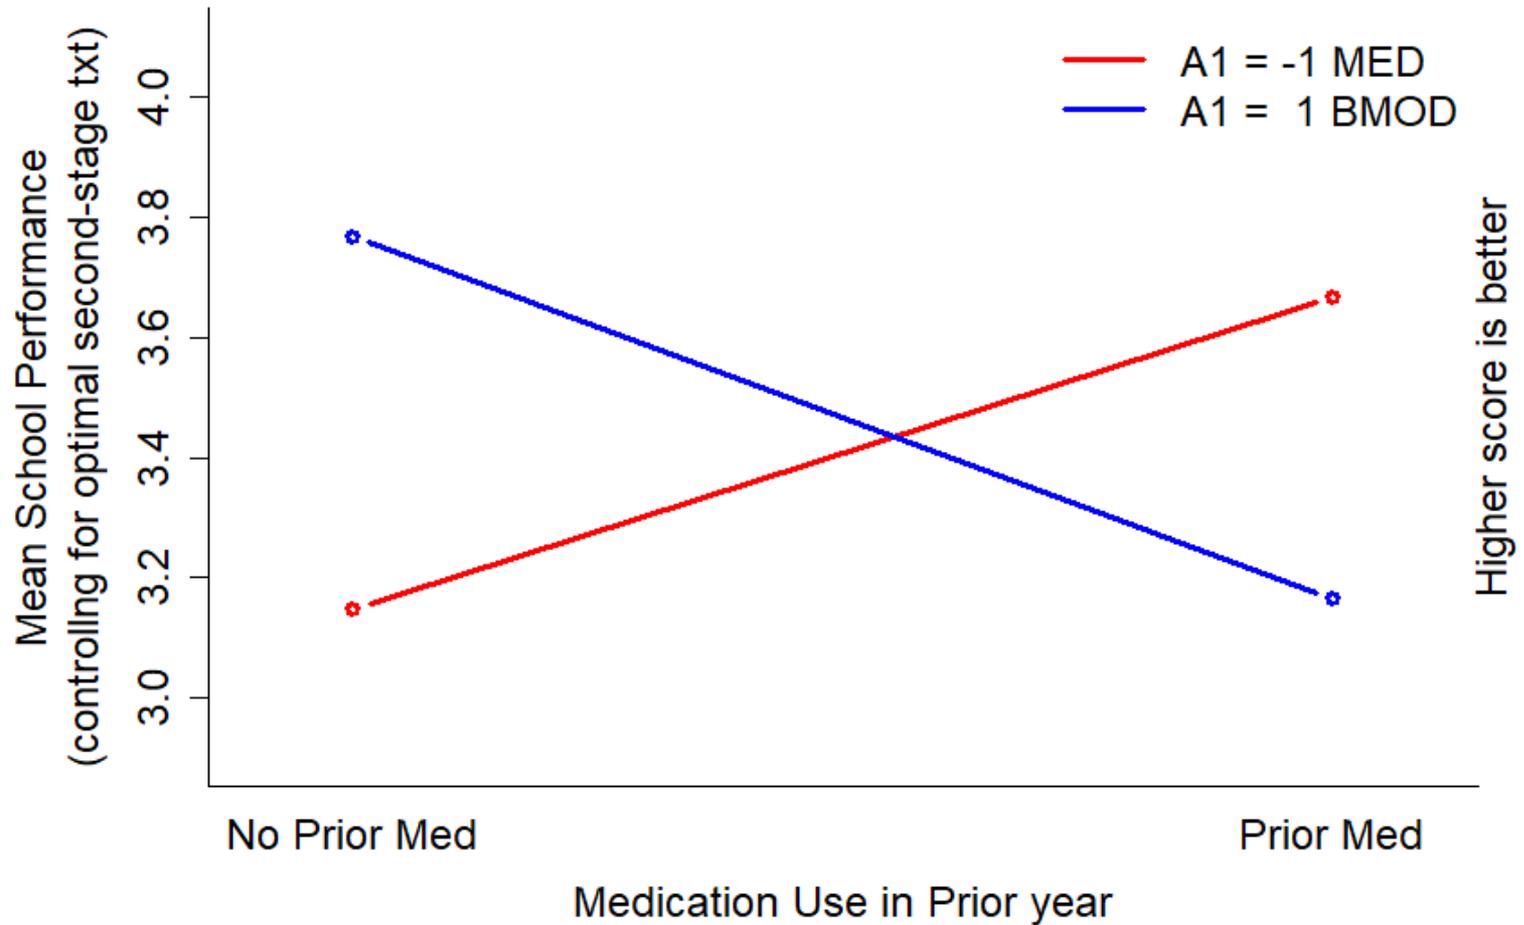
$$E[\hat{Y} \mid O_1, A_1] = \beta_0 + \beta_1 O_{11c} + \beta_2 O_{12c} + \beta_3 O_{14c} + \beta_4 \text{priorMed} + \beta_5 A_1 + \beta_6 (\text{priorMed} \times A_1)$$

Controlling for stage 2 tactic

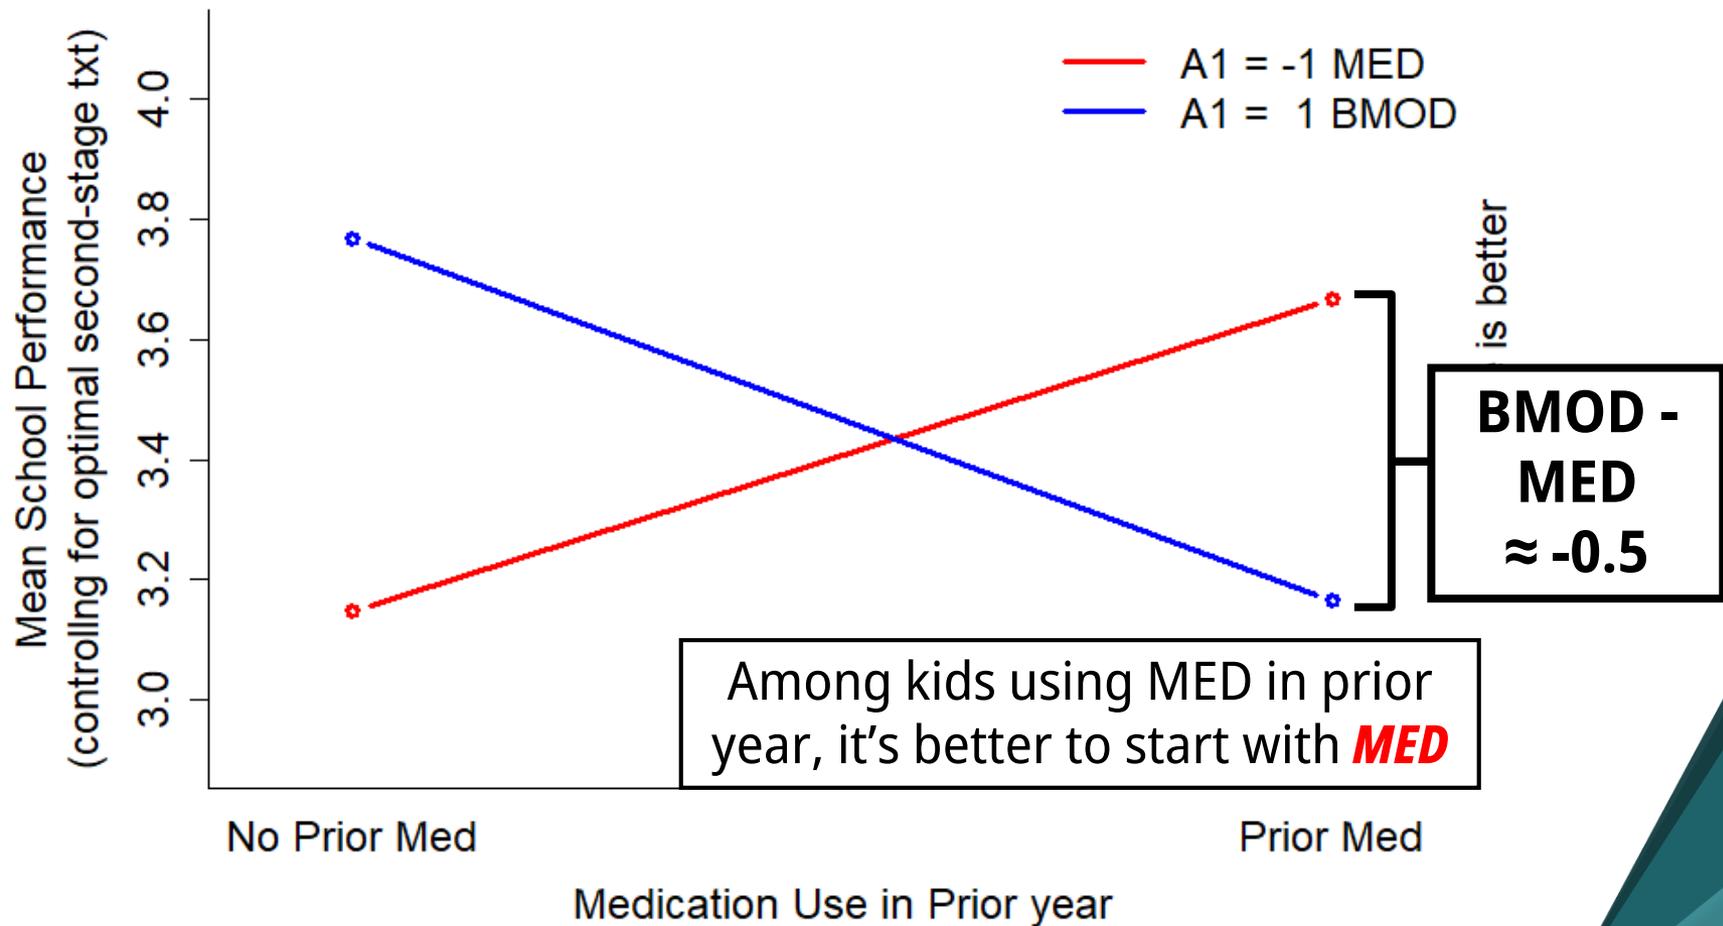
***This model will help us to:***

- Determine whether the best first stage option varies depending on whether the child received medication in prior year; and*
- Identify the best first stage option for children who received med in prior year vs. those that did not.*

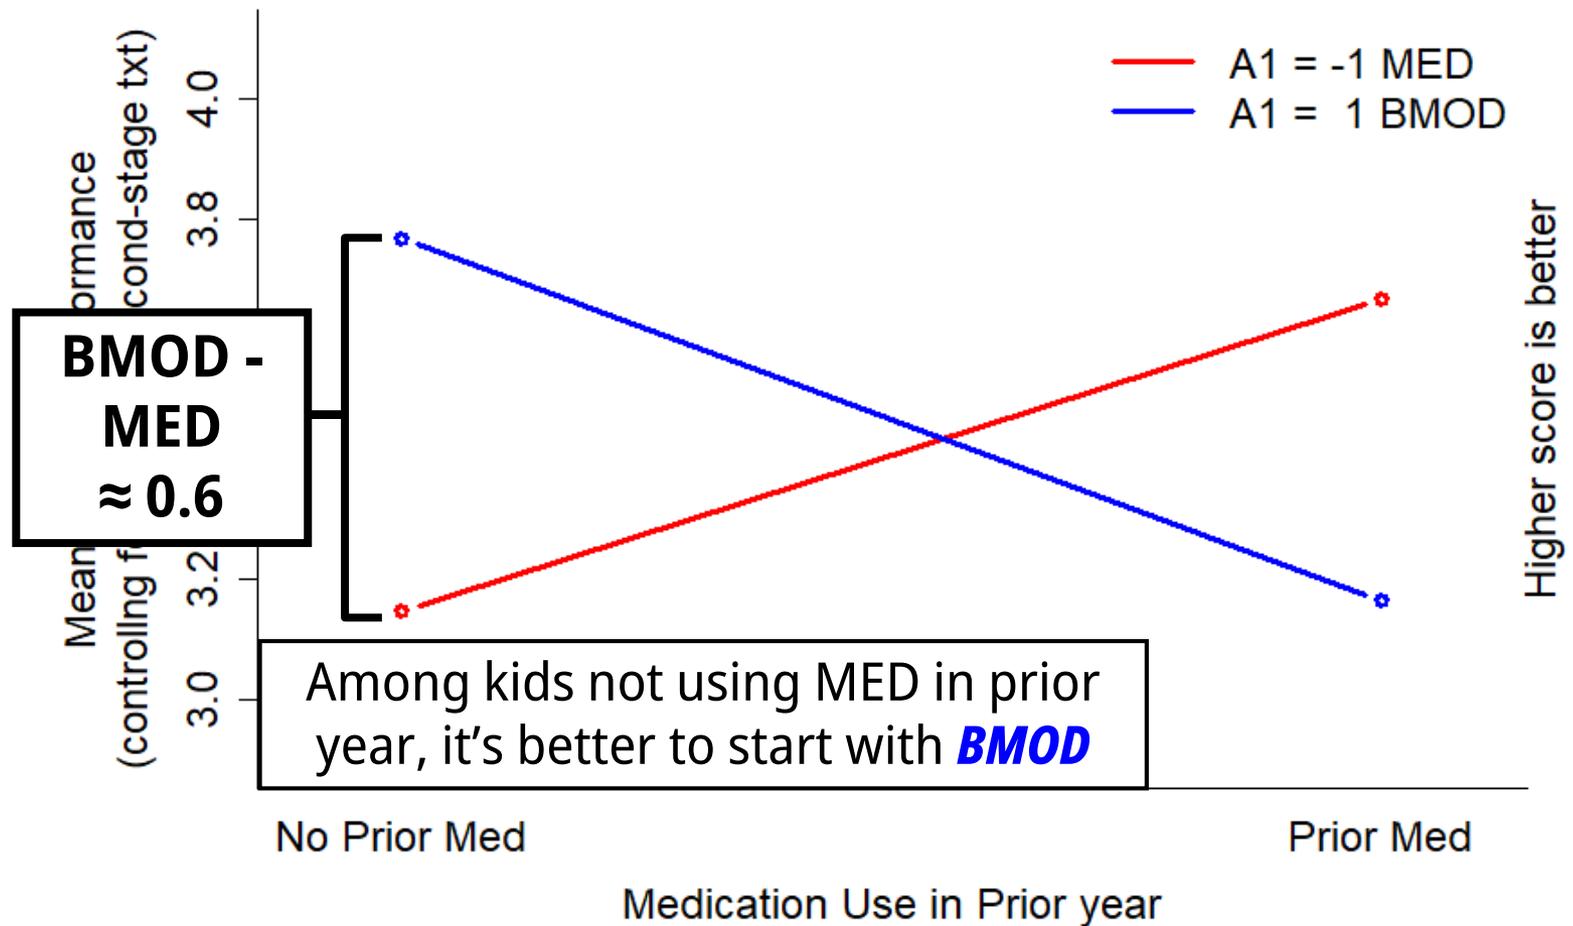
# Step 3: Move backwards to first-stage tailoring



# Step 3: Move backwards to first-stage tailoring



# Step 3: Move backwards to first-stage tailoring



# Summary of Q-learning results

## Step 1

- ***Second-Stage Regression***
  - Are **O1, A1, and O2** useful in making decisions about second-stage tactics?
  - (Are **O1, A1, and O2** useful in deciding **which NR** would benefit from Augment vs. Intensify?)

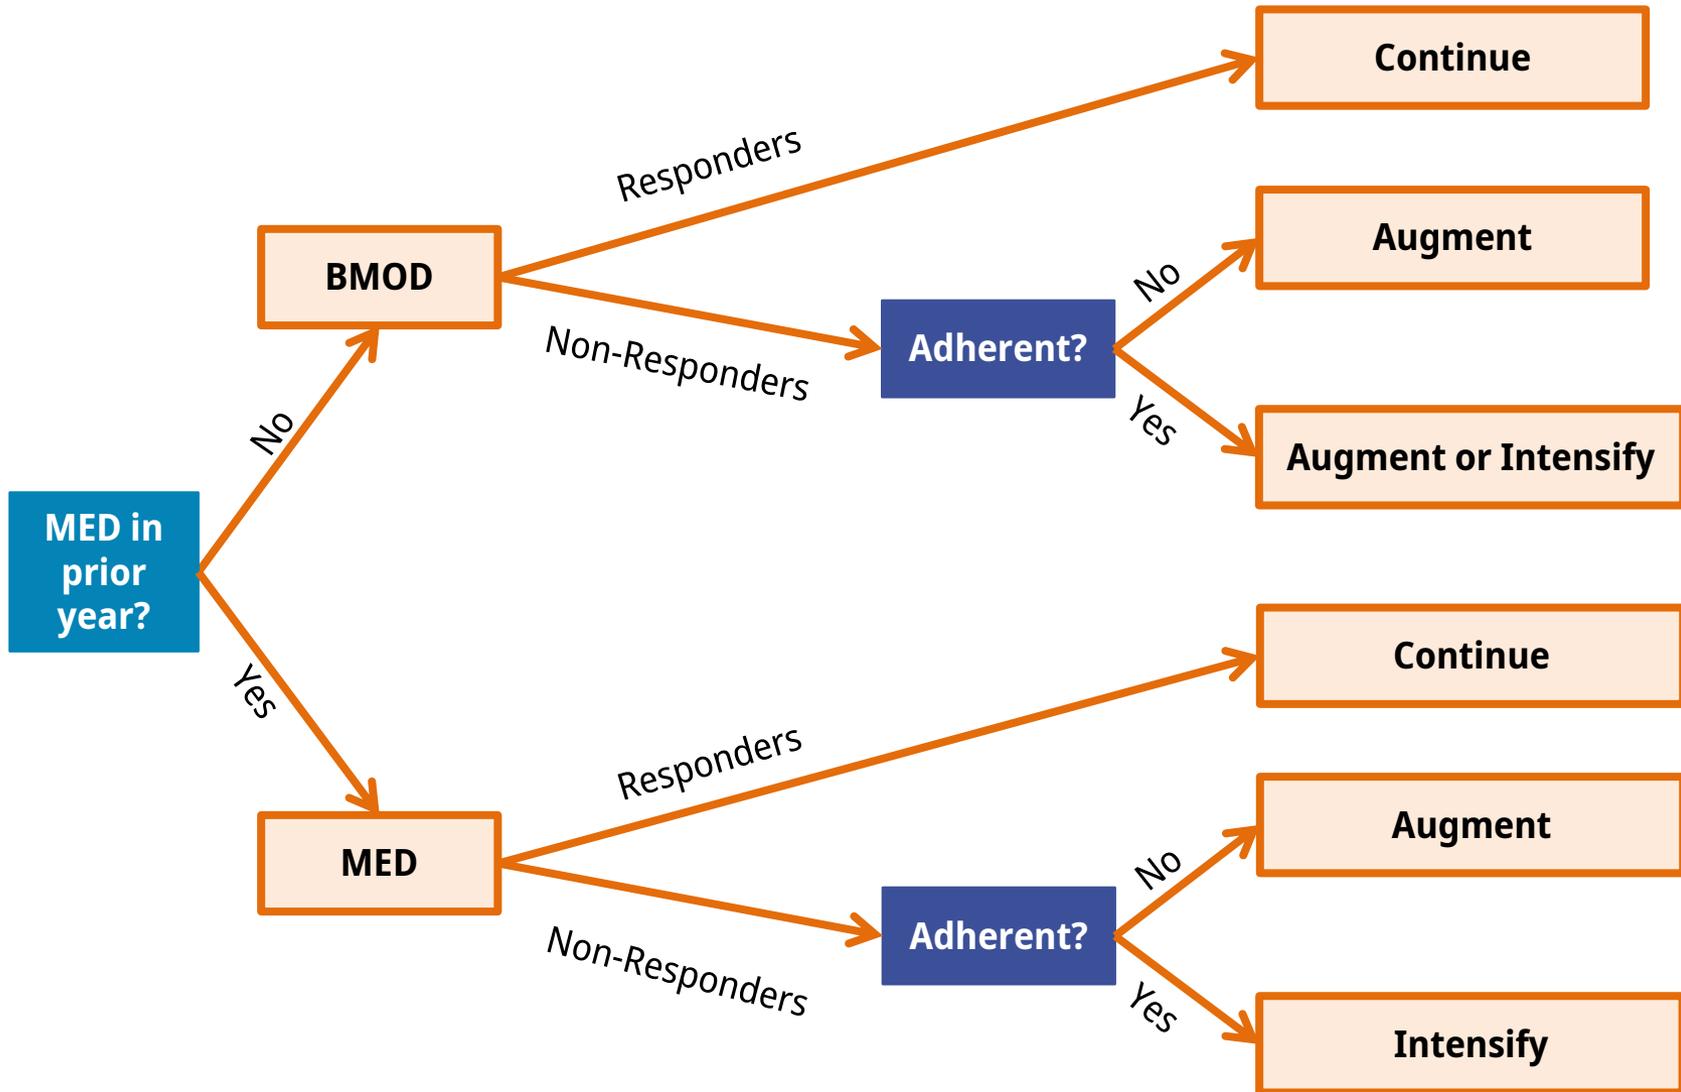
## Step 2

- ***Calculate  $\hat{Y}_i$*** 
  - What would the outcome be if they had received the best second-stage tactic given **O1, A1, and O2**?
  - $\hat{Y}_i$  is the **estimated optimal outcome under the best second-stage tactic for non-responders**. ( $\hat{Y}_i=Y$  for responders)

## Step 3

- ***First-Stage Regression***
  - Is **O1** useful in making decisions about first-stage tactics, *assuming we use optimal second-stage tactic*? (Use  $\hat{Y}_i$  from Step 2 as the outcome!)
  - (Is **O1** useful in deciding who would benefit from MED vs. BMOD, *assuming NRs get the best second-stage treatment*?)

# The estimated more-deeply tailored AI is



# 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 BMOD = {YES},
      THEN Stage 2 = {AUGMENT} or {INTENSIFY}.
  ELSE IF response status to Stage 1 = {R}
    THEN CONTINUE Stage 1.
```

# SAS code for 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 code for Q-learning

## Second-Stage Regression:

$$E[Y | \dots O_2, A_2] = \beta_0 + \beta_1 O_{11c} + \beta_2 O_{12c} + \beta_3 O_{13c} + \beta_4 O_{14c} + \beta_5 O_{12c} \\ + \beta_6 O_{21c} + \beta_7 \mathbf{A}_1 + \beta_8 \mathbf{O}_{22} \\ + \beta_9 A_2 + \beta_{10} (A_2 \times \mathbf{A}_1) + \beta_{11} (A_2 \times \mathbf{O}_{22})$$

\* Use only non-responders;

```
data dat10; set dat1; if R=0; run;
```

```
proc genmod data = dat10;
```

```
model y = o11cnr o12cnr o13cnr o14cnr o21cnr a1 o22 a2 a2*a1 a2*o22;
```

```
* INTENSIFY vs. AUGMENT when Stage 1 = MED, by ADHERENCE status;
```

```
estimate 'Diff: INT vs AUG for NR ADH MED' a2 2 a2*a1 -2 a2*o22 2 ;
```

```
estimate 'Diff: INT vs AUG for NR Non-ADH MED' a2 2 a2*a1 -2 a2*o22 0 ;
```

```
* INTENSIFY vs. AUGMENT when Stage 1 = BMOD, by ADHERENCE status;
```

```
estimate 'Diff: INT vs AUG for NR ADH BMOD' a2 2 a2*a1 2 a2*o22 2 ;
```

```
estimate 'Diff: INT vs AUG for NR Non-ADH BMOD' a2 2 a2*a1 2 a2*o22 0 ;
```

```
run;
```

# SAS code for Q-learning

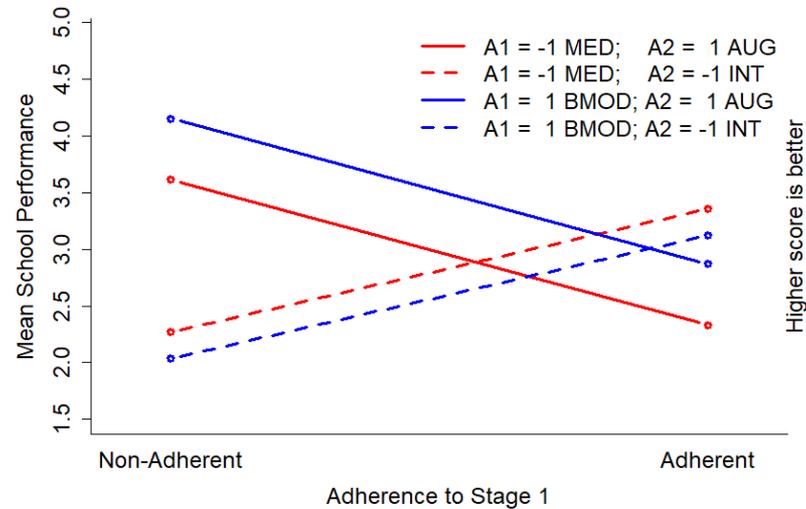
## Second-Stage Regression:

$$E[Y | \dots O_2, A_2] = \beta_0 + \beta_1 O_{11c} + \beta_2 O_{12c} + \beta_3 O_{13c} + \beta_4 O_{14c} + \beta_5 O_{12c} \\ + \beta_6 O_{21c} + \beta_7 A_1 + \beta_8 O_{22} \\ + \beta_9 A_2 + \beta_{10}(A_2 \times A_1) + \beta_{11}(A_2 \times O_{22})$$

### Contrast Estimate Results

Label	Mean Estimate	95% Confidence Limits		Standard Error	Pr > ChiSq
		Lower	Upper		
Diff: INT vs AUG for NR ADH MED	<b>1.0240</b>	0.4131	1.6350	0.3117	<b>0.0010</b>
Diff: INT vs AUG for NR Non-ADH MED	<b>-1.3412</b>	-1.9896	-0.6927	0.3308	<b>&lt;.0001</b>
Diff: INT vs AUG for NR ADH BMOD	<b>0.2503</b>	-0.3950	0.8956	0.3292	<b>0.4471</b>
Diff: INT vs AUG for NR Non-ADH BMOD	<b>-2.1149</b>	-2.7050	-1.5248	0.3011	<b>&lt;.0001</b>

# Step 1 results of Q-learning



## Contrast Estimate Results

Label	Mean Estimate	95% Confidence Limits		Standard Error	Pr > ChiSq
		Lower	Upper		
Diff: INT vs AUG for NR ADH MED	<b>1.0240</b>	0.4131	1.6350	0.3117	<b>0.0010</b>
Diff: INT vs AUG for NR Non-ADH MED	<b>-1.3412</b>	-1.9896	-0.6927	0.3308	<b>&lt;.0001</b>
Diff: INT vs AUG for NR ADH BMOD	<b>0.2503</b>	-0.3950	0.8956	0.3292	<b>0.4471</b>
Diff: INT vs AUG for NR Non-ADH BMOD	<b>-2.1149</b>	-2.7050	-1.5248	0.3011	<b>&lt;.0001</b>

# 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) button

# What does PROC QLEARN provide?

1. Data set with O1, A1, R, O2, A2, Y
2. The first regression model (best Stage 2 tactic)

$$Y \sim O1, A1, O2, A2$$

*NB: Be sure to specify sub-sample for this regression  
(non-responders in ADHD SMART)*

3. The second regression model (best Stage 1 option)

$$\hat{Y} \sim O1, A1$$

# What does PROC QLEARN do?

- 1. Implements Step 1:** Regression 1 for Stage 2 Tactic
  - Provides regression parameter estimates for determining best Stage 2 tactic
- 2. Implements Step 2:** Estimates  $\hat{Y}$ 
  - Assigns non-responders the outcome under the best Stage 2 tactic from Step 1.
  - Assigns responders their observed outcome.
- 3. Implements Step 3:** Regression 2 for Stage 1 Option
  - Provides regression parameter estimates for determining best Stage 1 option
  - Provides appropriate confidence intervals for the Stage 1 estimates

# Code for 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 regression  
RUN;
```

### ***Regression 1 (Stage 2):***

$INT + MAIN2 + TAILOR2 + TAILOR2 * STG2TRT + STG2TRT$

### ***Regression 2 (Stage 1):***

$INT + MAIN1 + TAILOR1 + TAILOR1 * STG1TRT + STG1TRT$

# Code for PROC QLEARN

```
DATA dat11; SET dat1; S = 1-R; RUN;
```

```
PROC QLEARN data=dat11 contrasts1=contrasts1 deriveci;  
  MAIN1 o11c o12c o14c;  
  TAILOR1 o13;  
  MAIN2 o11c o12c o13c o14c o21c;  
  TAILOR2 a1 o22;  
  STG2SAMPLE s;  
  RESPONSE y;  
  STG1TRT a1;  
  STG2TRT a2;  
RUN;
```

Request contrasts of interest ("estimates" in GENMOD). See next slide...

Ask for confidence intervals by Laber and Murphy (2012)

*This will fit two regressions:*

$$\begin{aligned} \text{Stage 2 (for NR's): } E[Y | A, O] = & \beta_0 + \beta_1 O_{11c} + \beta_2 O_{12c} + \beta_3 O_{13c} + \beta_4 O_{14c} \\ & + \beta_5 O_{12c} + \beta_6 O_{21c} + \beta_7 A_1 + \beta_8 O_{22} \\ & + \beta_9 A_2 + \beta_{10} (A_2 \times A_1) + \beta_{11} (A_2 \times O_{22}) \end{aligned}$$

$$\begin{aligned} \text{Stage 1: } E[\hat{Y} | O, A] = & \beta_0 + \beta_1 O_{11c} + \beta_2 O_{12c} + \beta_3 O_{14c} \\ & + \beta_4 O_{13} + \beta_5 A_1 + \beta_6 (O_{13} \times A_1) \end{aligned}$$

# Code for 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 011c + b3 012c + b4 014c + b5 013 + b6 A1*013 + 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;
```

# PROC QLEARN Results

## First Stage Regression Result

Variable	Parameter Estimates	Confidence Upper	Confidence Lower	Interval
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	
<b>o13</b> <b>:a1</b>	<b>-0.5610</b>	<b>-0.2836</b>	<b>-0.8292</b>	
a1	0.3104	0.4992	0.0993	
Contrasts	Parameter Estimates	Confidence Upper	Confidence Lower	Interval
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	

# PROC QLEARN Results

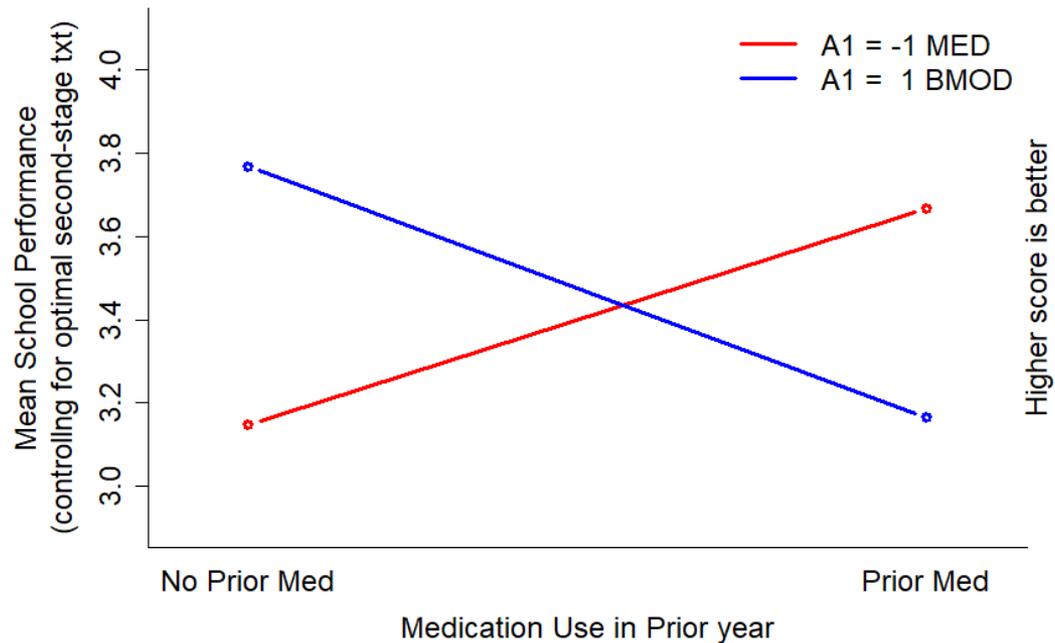
## First Stage Regression Result

---

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

mean under BMOD for kids w prior med  
mean under MED for kids w prior med  
mean diff (BMOD-MED) kids w prior med  
mean under BMOD for kids w no prior med  
mean under MED for kids w no prior med  
mean diff (BMOD - MED) kids w no prior med

# PROC QLEARN Step 3 Results



Contrasts	Parameter Estimates	Confidence Interval	
		Upper	Lower
(BMOD-MED) prior med	-0.5012	-0.0032	-1.0399
(BMOD - MED) no prior med	0.6208	0.9984	0.1986

# 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) button

# What we learned 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 we learned 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 BMOD = {YES},
      THEN Stage 2 = {AUGMENT} or {INTENSIFY}.
  ELSE IF response status to Stage 1 = {R}
    THEN CONTINUE Stage 1.
```

# What we learned from Q-learning

The **mean Y**, school performance, under the more deeply tailored AI obtained via Q-learning is estimated to be **3.72**.

- Recall (BMOD, AUGMENT) was the AI with the largest mean among the 4 embedded AIs (only tailored on response)
- The value of the more deeply-tailored AI is larger than the value of the AI that started with BMOD and used AUGMENT for non-responders (mean = **3.51**)

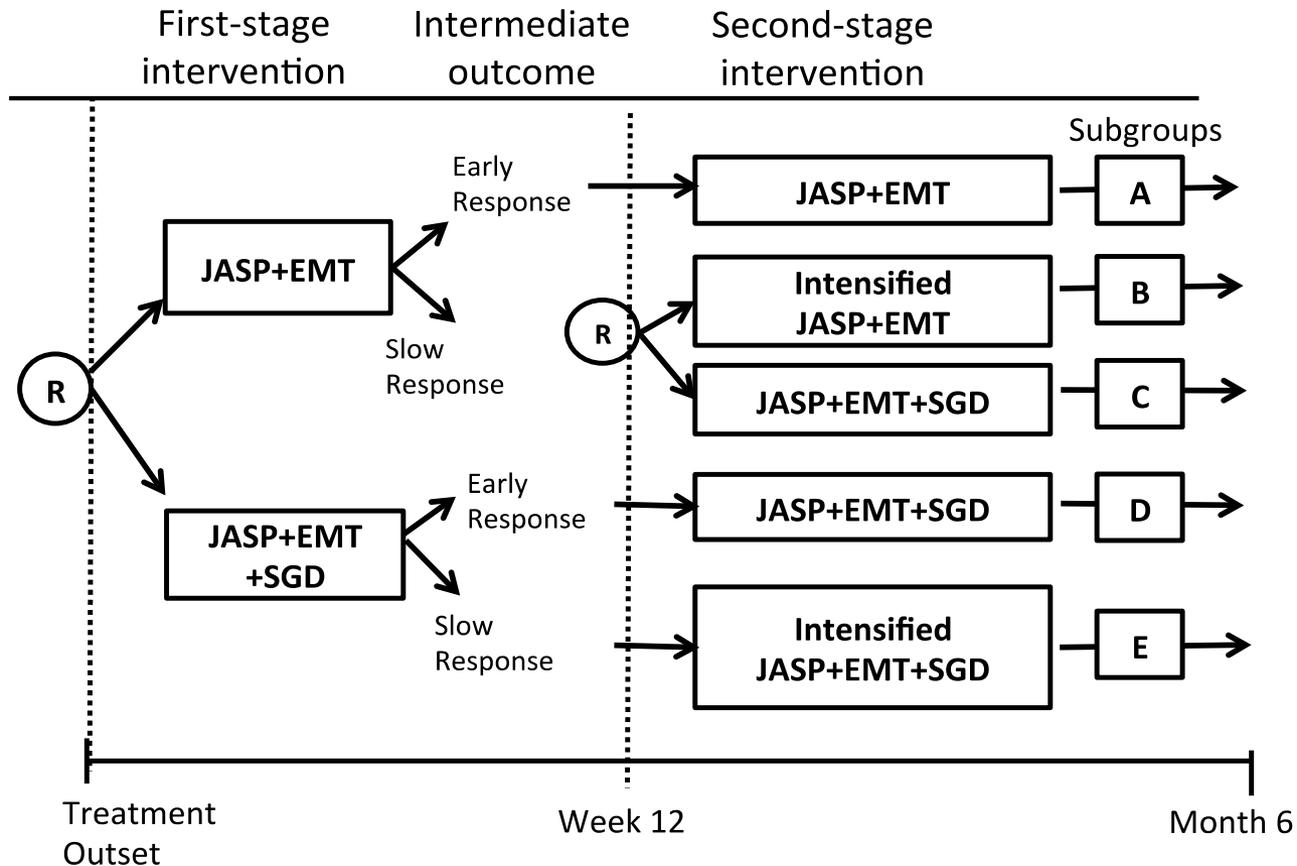
# What we learned from Q-learning

We may want to evaluate the efficacy of this *proposed AI* versus a suitable control (e.g., usual care) with a subsequent trial (i.e., RCT)

# 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: Autism SMART



**JASP** → Joint Attention and joint Engagement

**EMT** → Enhanced Milieu Teaching

**SGD** → Speech Generating Device (e.g., an iPad or DynaVox)

# “Manual” Q-learning for the interested

- 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

# “Manual” Q-learning: Step 2

```
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

# “Manual” Q-learning: Step 3

\* 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

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