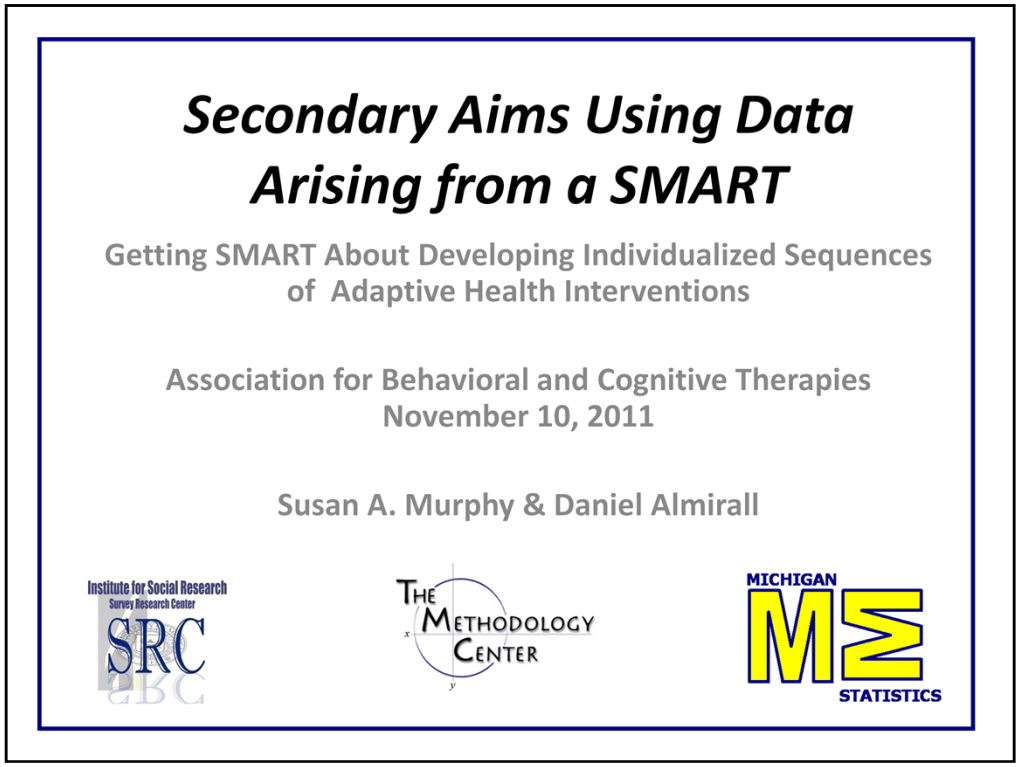

50 minutes

This is the 5th (final) module of a 5-module Seminar on experimental designs for building optimal adaptive health interventions.

By now, you know what an ATS is. You have discussed why they are important in terms of managing chronic disorders (indeed, an ATS 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 the data analysis methods used to address typical secondary research aims posed in SMART trials.

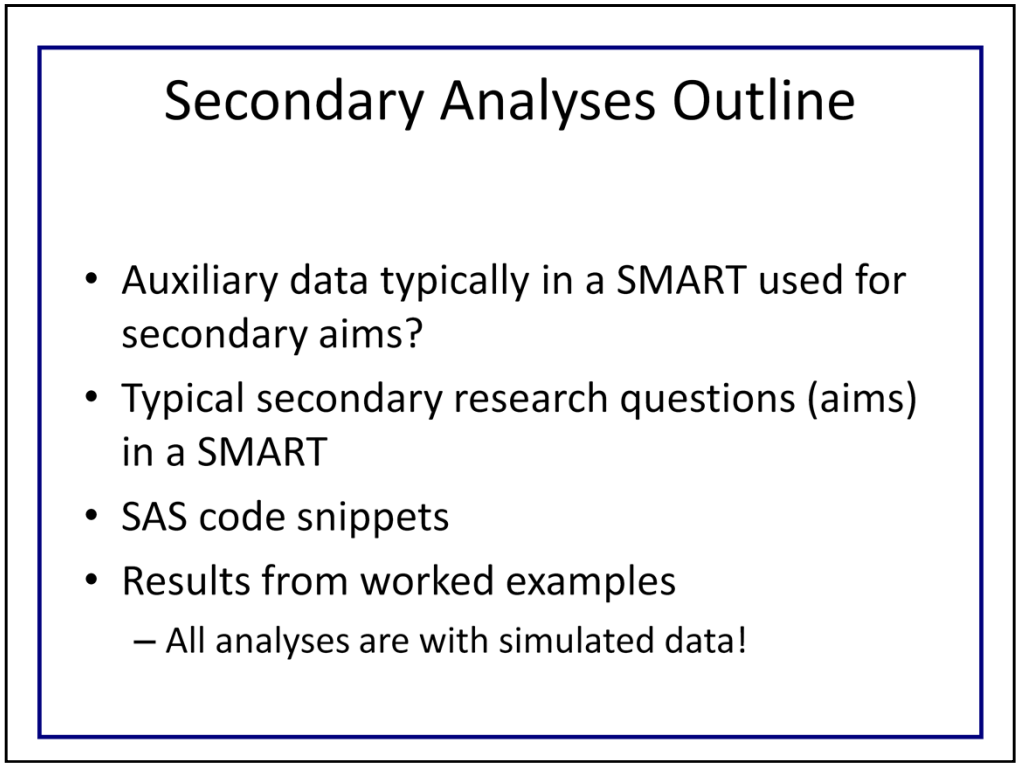

Outline

The basic idea is I will describe a type of question using words, I will describe it using a picture, then I will show you SAS code to answer that question, and results from a worked example using simulated data (that you can practice with later).

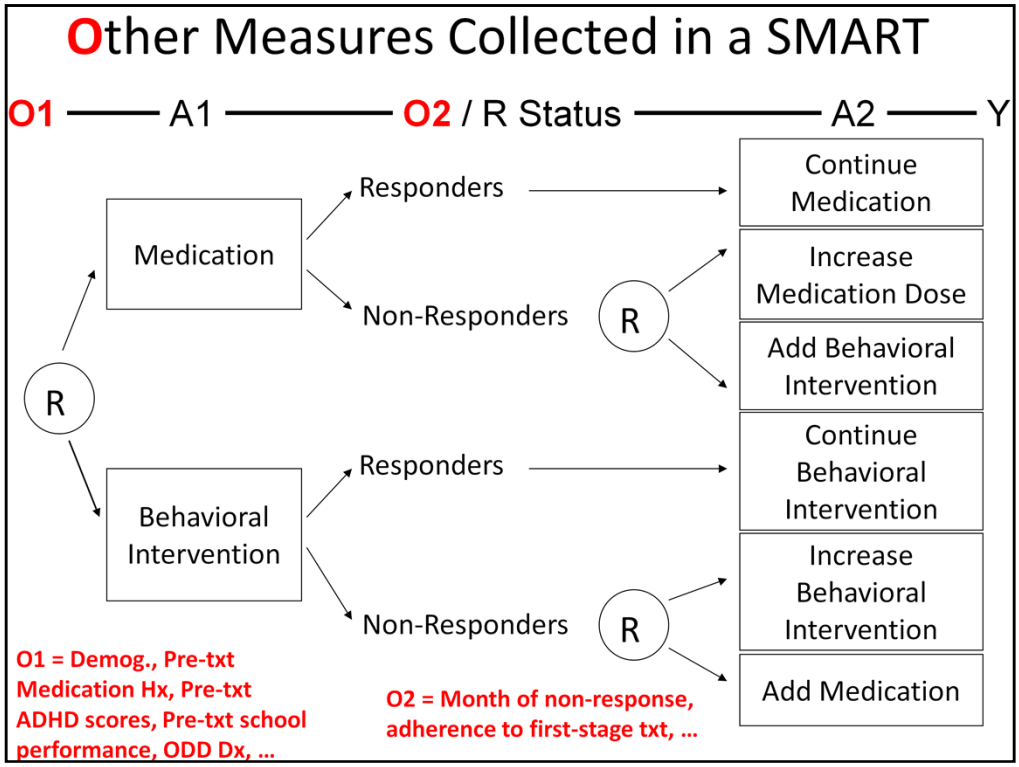

In addition to standard outcomes scales/measures, many other things could me measured during initial treatment (in this SMART study) that could be used in secondary analyses to more deeply tailor subsequent treatment, including:

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

etc…

Typical Secondary Aim 1: Best second-line tactic?

- Among children who do not respond to (either) first-line treatment, is it better to increase initial treatment or to add a different treatment to the initial treatment?
  - Regardless of history of treatment.

[explained with a picture on the next slide….]

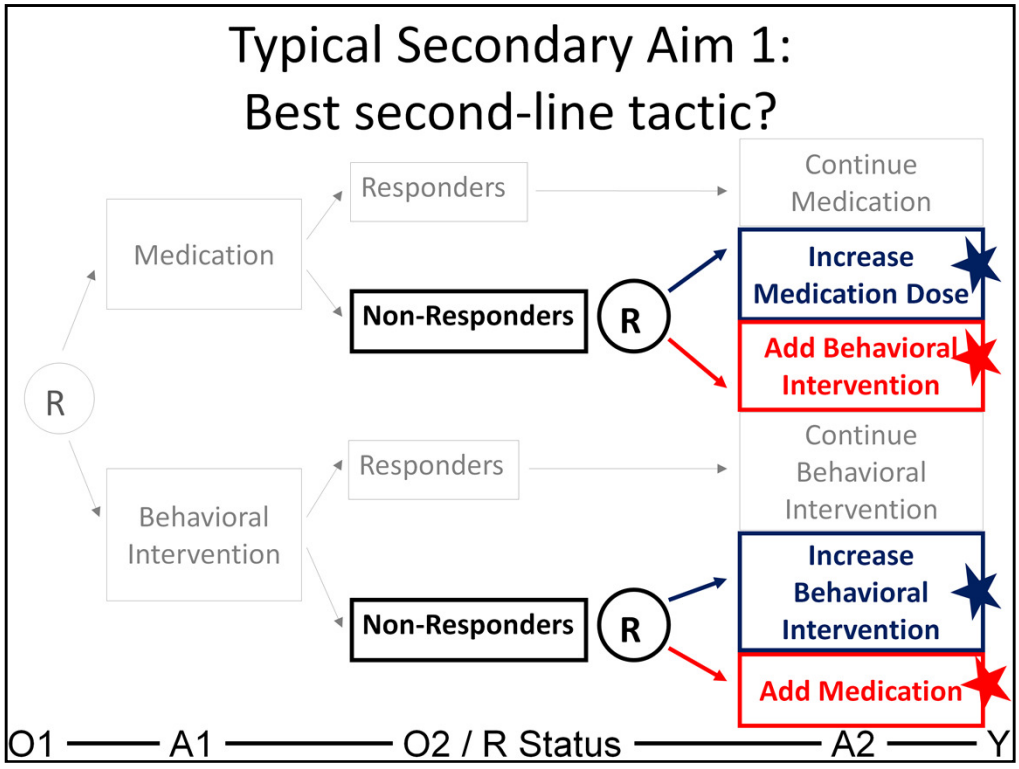

This is not a comparison of adaptive treatment strategies, per se. Rather it informs the tactical decision often made in clinical practice of whether to add to the treatment with something new versus increase the dosage/intensity of treatment.

Note that this is a comparison of the blue star cells versus the red star cells, pooled over (averaged over) first-line. The pooling leads to more power (i.e., larger sample size for the comparison of tactics) but the pooling does not always make sense. Here it does if we think of it from a mental health services delivery point of view.

SAS code and results for Secondary Aim 1: Second-line tactic

```
* use only non-responders;
data dat4;
  set dat1; if R=0;
run;
* simple comparison to compare mean Y on add vs intensify (A2);
proc genmod data = dat4;
  model y = a2 o11c o12c o13c;
  estimate 'Mean Y w/INTENSIFY tactic' intercept 1 a2  1;
  estimate 'Mean Y w/ADD TXT tactic'   intercept 1 a2 -1;
  estimate 'Between groups difference'              a2  2;
run;
```

Contrast Estimate Results

		95% Conf Limits		
Label	**Estimate**	**Lower**	**Upper**	**P-value**
Mean Y w/INTENSIFY tactic	3.2143	2.9026	3.5260	<.0001
Mean Y w/ADD TXT tactic	3.4255	3.1308	3.7202	<.0001
Between groups difference	-0.2112	-0.6402	0.2177	0.3345

This analysis is with simulated data.

On average, the tactic of ADDING is better, but not statistically significantly better.

We can also break down the effect of second-line treatment depending on initial treatment status. We do this next…

Typical Secondary Aim 2: Best second-line treatment?

a. Among children who do not respond to first-line medication, is it better to increase dosage or to add behavioral modification?

b. Among children who do not respond to first-line behavioral modification, is it better to increase intensity of behavioral treatment or to add medication?

[explained with a picture on the next slide….]

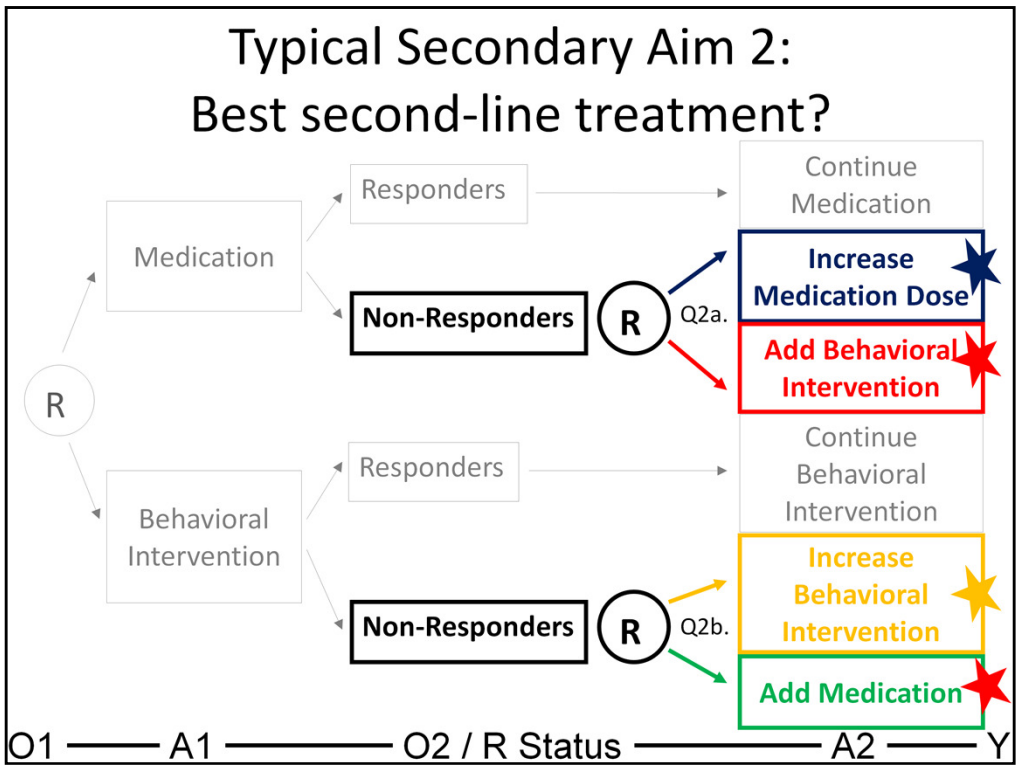
Typical Secondary Aim 2:
Best second-line treatment?
R
Medication
Responders
Non-Responders
R
Q2a.
Continue Medication
Increase Medication Dose
Add Behavioral Intervention
Behavioral Intervention
Responders
Non-Responders
R
Q2b.
Continue Behavioral Intervention
Increase Behavioral Intervention
Add Medication
O1
A1
O2 / R Status
A2
Y

SAS code and results for Secondary Aim 2a: Second-line txt after MED

```
* use only medication non-responders;
data dat2;
  set dat1; if R=0 and A1=-1;
run;
* simple comparison to compare mean Y on add vs intensify (A2);
proc genmod data = dat2;
  model y = a2     ;
  estimate 'Mean Y w/INTENSIFY MED' intercept 1 a2  1;
  estimate 'Mean Y w/ADD BMOD'      intercept 1 a2 -1;
  estimate 'Between groups difference'          a2  2;
run;
```

Contrast Estimate Results

		95% Conf Limits		
Label	**Estimate**	**Lower**	**Upper**	**P-value**
Mean Y w/INTENSIFY MED	3.5714	3.0862	4.0567	<.0001
Mean Y w/ADD BMOD	3.2500	2.8440	3.6560	<.0001
Between groups difference	0.3214	-0.3113	0.9541	0.3194

This analysis is with simulated data.

Slightly positive, but No statistically significant difference between INT vs ADD, among non-responders to MED.

SAS code and results for Secondary Aim 2b: Second-line txt after BMOD

```
* use only BMOD non-responders;
data dat3;
  set dat1; if R=0 and A1=1;
run;
* simple comparison to compare mean Y on add vs intensify (A2);
proc genmod data = dat3;
  model y = a2 o11c o12c o13c;
  estimate 'Mean Y w/INTENSIFY BMOD' intercept 1 a2  1;
  estimate 'Mean Y w/ADD MED'        intercept 1 a2 -1;
  estimate 'Between groups difference'           a2  2;
run;
```

Contrast Estimate Results

		95% Conf Limits		
Label	**Estimate**	**Lower**	**Upper**	**P-value**
Mean Y w/INTENSIFY BMOD	3.0357	2.6436	3.4278	<.0001
Mean Y w/ADD MED	3.5556	3.1563	3.9548	<.0001
Between groups difference	-0.5198	-1.0795	0.0398	0.0687

This analysis is with simulated data.

Among non-responders to BMOD, it is better to ADD MED rather than to INTENSIFY BMOD.

Typical Secondary Aim 3: Second-line treatment tailoring?

a. Does adherence to first-line MED strongly moderate the impact of increasing MED dosage versus adding BMOD?

b. Does adherence to first-line BMOD strongly moderate the impact of intensifying BMOD versus adding MED?

By “strongly moderate” we mean such that adherence can be used as a tailoring variable in a decision rule for whether to intensify versus add txt.

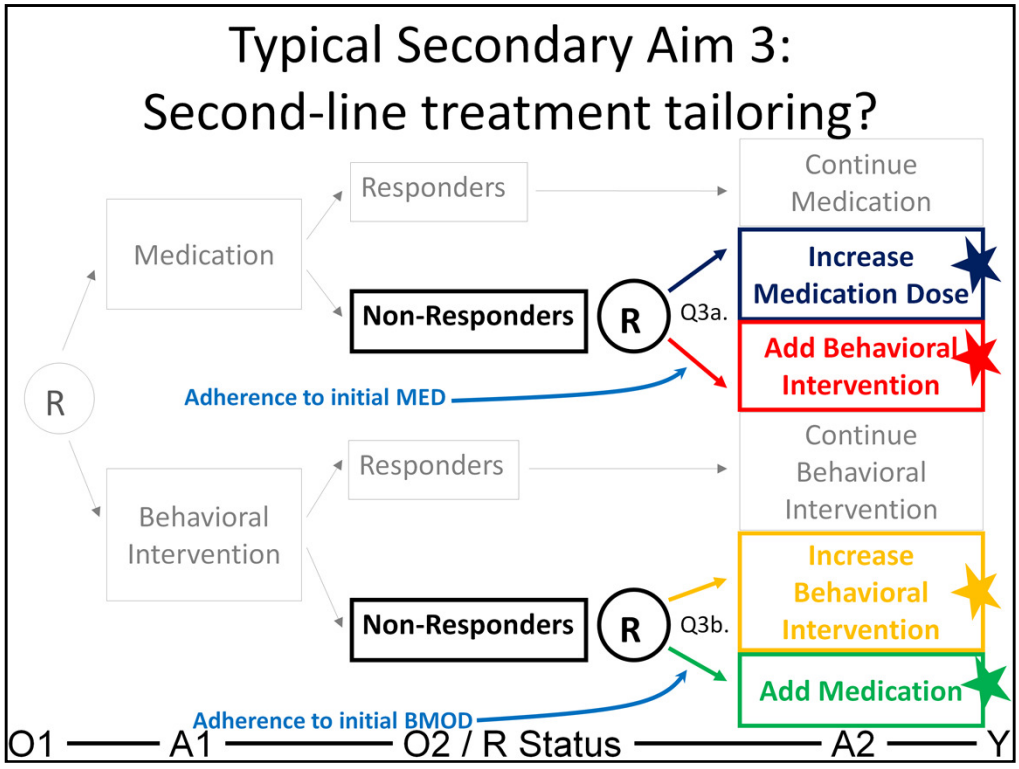

Why is this such an interesting question? Because if adherence to initial treatment strongly moderates the impact of increase/add on Y, then it can serve as a tailoring variable (or one that we investigate in more detail in the next randomized trial).

SAS code and results for Secondary Aim 3: Second-line treatment tailoring

```
* use only non-responders;
data dat5; set dat1; if R=0; run;

* comparison of add vs intensify given first line txt and adherence;
proc genmod data = dat5;
  model y = o11c o12c o13c a1 a1*o11c o21c o22 a2 a2*a1 a2*o22;
  * effect of add vs intensify given first-line = MED x ADH status;
  estimate 'INT vs ADD for NR MED ADH'     a2 2 a2*a1 -2 a2*o22  2 ;
  estimate 'INT vs ADD for NR MED Non-ADH' a2 2 a2*a1 -2 a2*o22  0 ;
  * effect of add vs intensify given first-line = BMOD x ADH status;
  estimate 'INT vs ADD for NR BMOD ADH'     a2 2 a2*a1 2 a2*o22  2 ;
  estimate 'INT vs ADD for NR BMOD Non-ADH' a2 2 a2*a1 2 a2*o22  0 ;
run;
```

Contrast Estimate Results

		95% Conf Limits		
Label	Estimate	Lower	Upper	P-value
INT vs ADD for NR MED ADH	1.0473	0.5682	1.5263	<.0001
INT vs ADD for NR MED Non-ADH	-1.5658	-2.1587	-0.9728	<.0001
INT vs ADD for NR BMOD ADH	1.2651	0.7529	1.7773	<.0001
INT vs ADD for NR BMOD Non-ADH	-1.3479	-1.7493	-0.9465	<.0001

This analysis is with simulated data.

Among non-responders who adhere to (either) first-line treatment, it is better to INTENSIFY treatment rather than ADD a different treatment (positive effects). Whereas, for non-responders who do not adhere to (either) first-line treatment, it is better to ADD than to INTENSIFY (negative effects).

These results + the fact that MED group had better adherence (see next slide) explains the previous results for best second line txt. We saw that for BMOD NR, it was better to ADD than INT; and for MED NR, it was better to INT than ADD. Well, this is bec the MED had better adherence, so more of the positive effect (1.0473) of INT was represented in MED group mean. Whereas BMOD group had slightly worse adherence, so that a more of the negative effect (-1.3479) was represented (and it was stronger in absolute terms anyway). To formalize this intuition the next slide shows the impact of first-line txt on ADH.

Side analysis: SAS code and results for impact of first-line treatment on ADH

```
proc freq data=dat1;
  table a1*o22 / chisq nocol nopercent;
run;
```

Frequency Row Pct	ADH = 0	ADH = 1	Total
A1 = -1 MED	28 41.18	40 58.82	68
A1 = 1 BMOD	52 63.41	30 36.59	82
	80	70	150

In terms of adherence, initial MED is better than initial BMOD by 22% (p-value < 0.01).

This analysis is with simulated data.

This confirms the explanation of the previous slide.

Let’s take a quick break!

Up next: A method for building a more deeply-tailored ATS.

Typical Secondary Aim 4: A more deeply individualized ATS via Q-learning

Q-Learning is an extension of regression to sequential treatments.

- Q-Learning results in a proposal for an adaptive treatment strategy with greater individualization.
- A subsequent trial would evaluate the proposed adaptive treatment strategy versus usual care.

This is likely the most interesting & fun aim that one can investigate using data from a SMART design.

The name “Q-learning” refers to learning more about the “Quality” of an adaptive treatment strategy.
This is an idea borrowed from computer scientists.

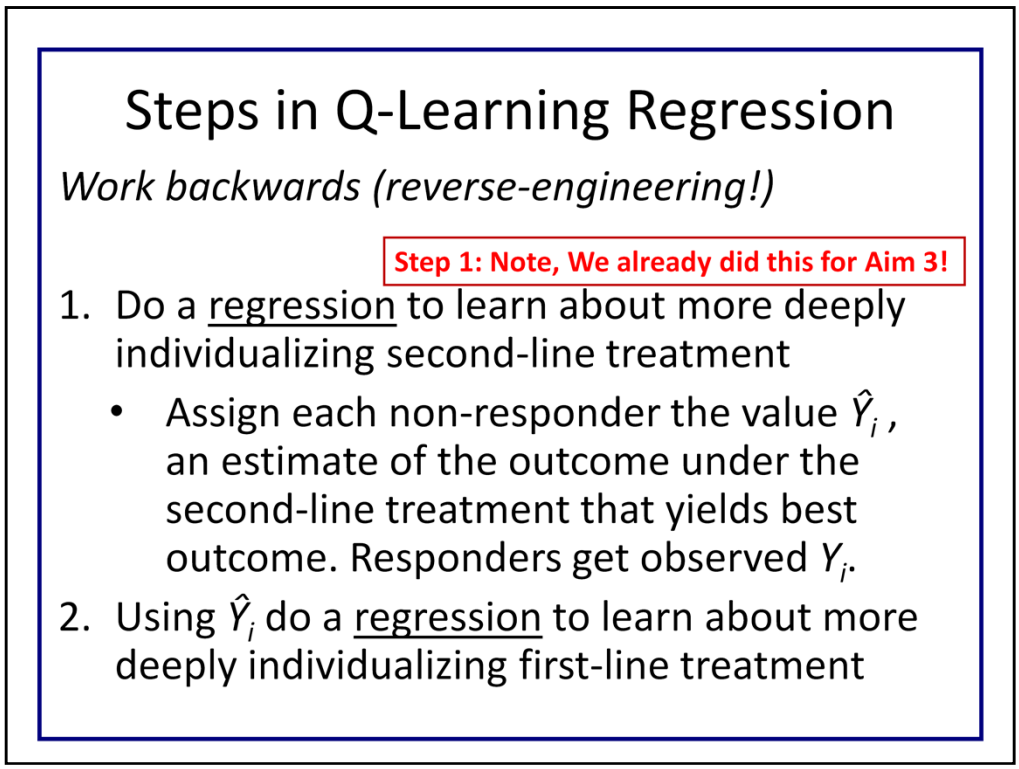

Steps:
First, do a regression at stage 2 to learn about the optimal second-line treatment given characteristics of the participant at baseline and outcome during first-line treatment

Second, do a regression using an outcome that already has taken into account future optimal treatment to learn about the optimal first-line treatment.

Conduct the regressions in backwards order! E.g. Stage 2 first, then Stage 1.
Why?

Stage 1 dependent variable must control for Stage 2 treatment.
Stage 1 dependent variable is a predictor of *Y* under optimal treatment in stage 2.
Stage 2 analysis is used to construct hat{Y}

We are going to demonstrate the Q-learning algorithm results with adherence as the candidate stage 2 tailoring variable
and the presence of oppositional defiant disorder ODD as the candidate stage 1 tailoring variable.
SAS code for doing qlearning is coming soon, but R code is available. In the package sent to you after this workshop, you'll receive the R code used to carry out these analyses. Due to lack of time, we will not go over the R code here.

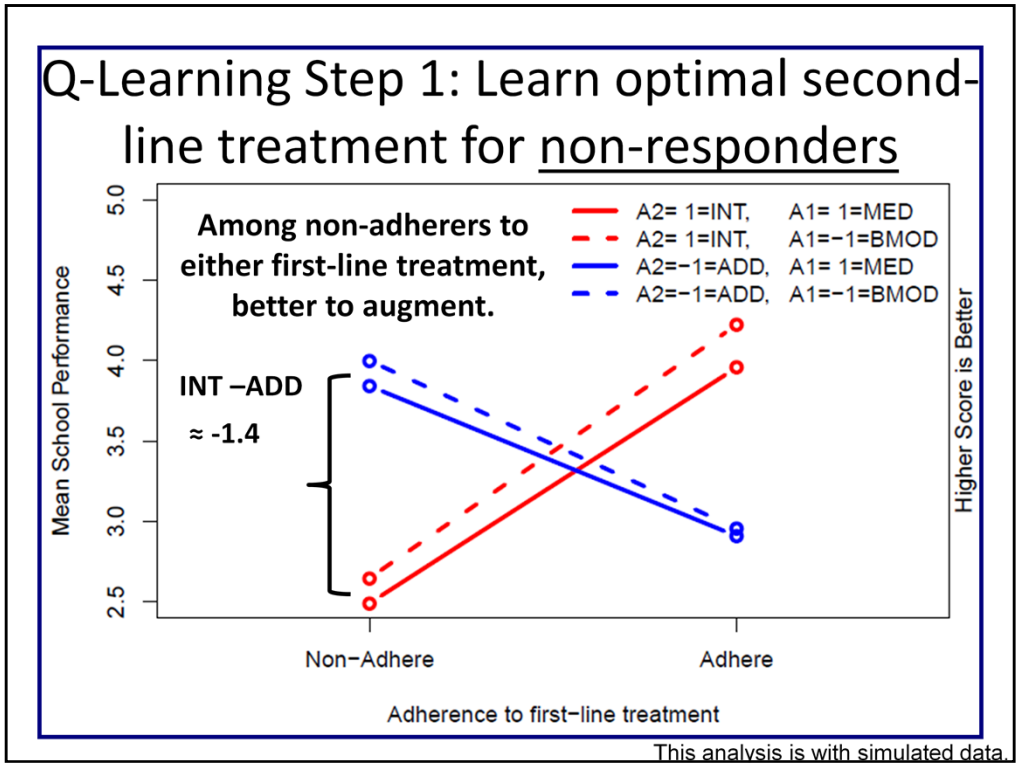

Quite large effect a la Cohen's effect sizes.

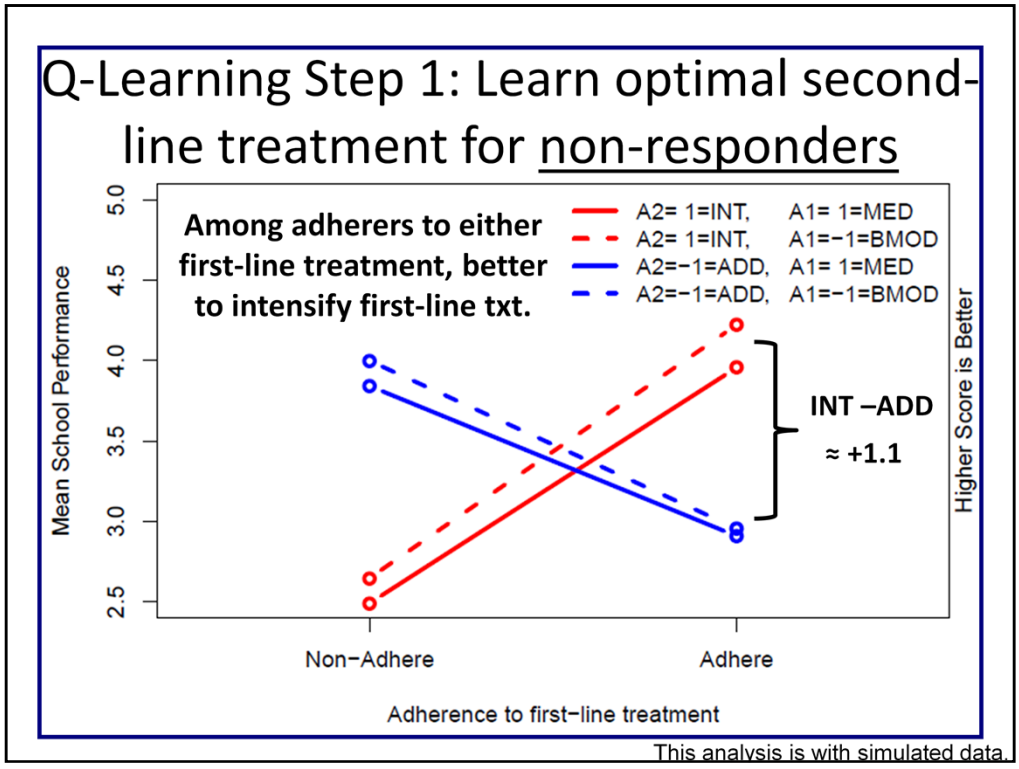

Large effect ala Cohen.

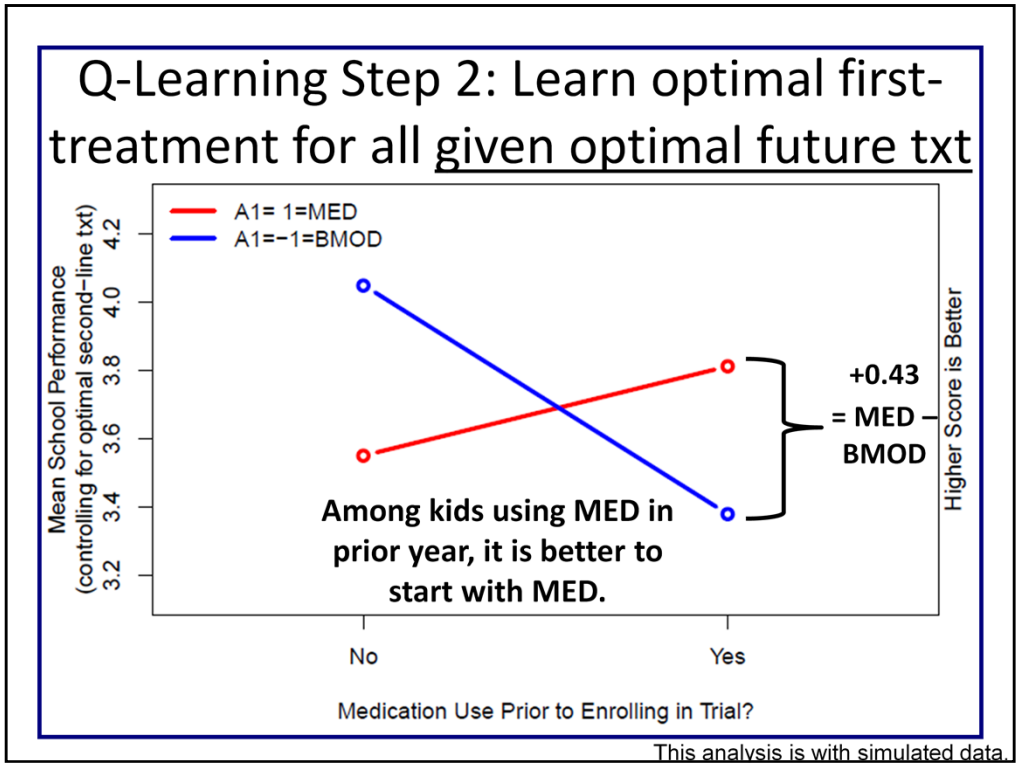
Q-Learning Step 2: Learn optimal first-treatment for all given optimal future txt
A1= 1=MED
A1=−1=BMOD
Mean School Performance (controlling for optimal second−line txt)
3.2
3.4
3.6
3.8
4.0
4.2
+0.43
= MED – BMOD
Higher Score is Better
Among kids using MED in prior year, it is better to start with MED.
No
Yes
Medication Use Prior to Enrolling in Trial?
This analysis is with simulated data.

So, we assign

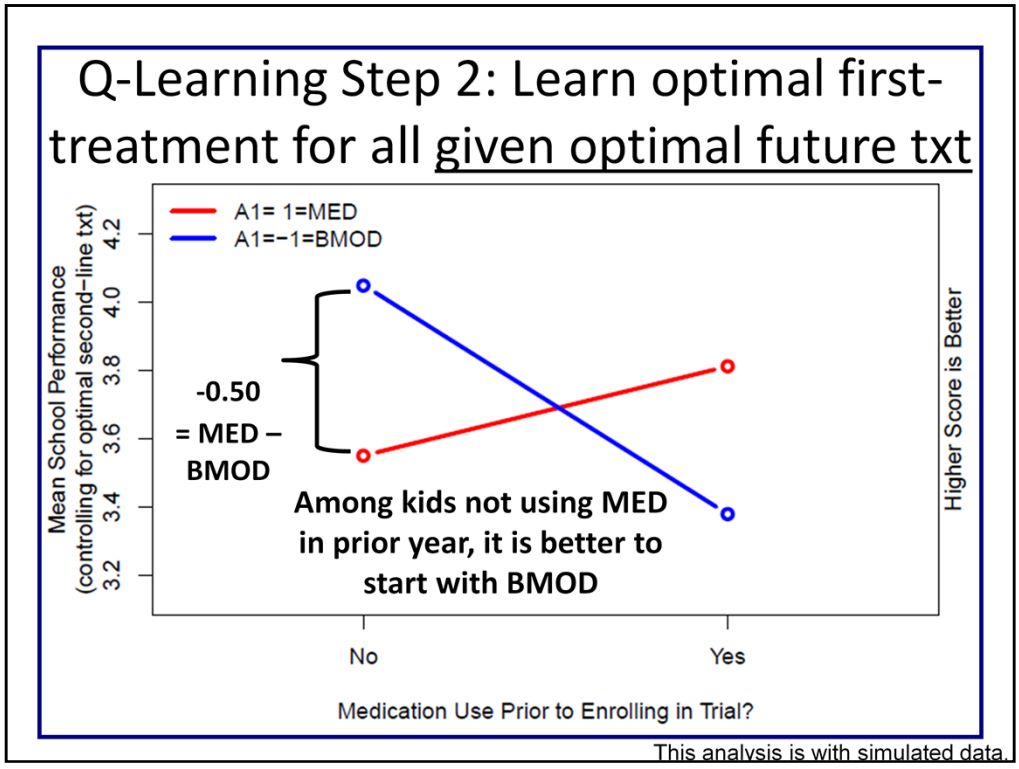

So, we assign

Both of these are moderately sized effects around ES=0.50 (“moderate” a la Cohen)

Note that this is not 1 of the 4 embedded ATS *as part of the trial design*. This is a more deeply individualized ATS that is a function of much more than just early response status. Indeed, this ATS is a function of prior MED use, first-line txt, response status, adherence to first-line txt, and second-line txt.

What did we learn with Q-learning?

Adaptive Treatment Strategy Proposal

- The mean Y, school performance, under the more deeply individualized ATS obtained via Q-learning is estimated to be 3.99.
- This is larger than the value of the ATS which started with BMOD and augmented with MED for non-responders (mean = 3.47)
  - (BMOD, MED) was the ATS with the largest mean among the 4 embedded ATSs.

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

Note that this is not 1 of the 4 embedded ATS *as part of the trial design*. This is a more deeply individualized ATS that is a function of much more than just early response status. Indeed, this ATS is a function of prior MED use, first-line txt, response status, adherence to first-line txt, and second-line txt.

Citations to Technical Reports

- Nahum-Shani, I., Qian, M., Almirall, D., Pelham, W. E., Gnagy, B., Fabiano, G., Waxmonsky, J., et al. (2010**).** ***Q-Learning: A data analysis method for constructing adaptive interventions*** **.** Technical Report, The Methodology Center, Penn State University.
- Nahum-Shani, I., Qian, M., Almirall, D., Pelham, W. E., Gnagy, B., Fabiano, G., Waxmonsky, J., Yu, J., & Murphy, S. (2010). ***Experimental design and primary data analysis for informing sequential decision making processes****.* Technical Report, The Methodology Center, Penn State University.

Practice Exercise

Exercise: Using the sample SMART you developed previously, write down a secondary research question of interest to you. What data analysis approach would you use to address this question?

Thank you.

- Software for Q-learning is now available in R and it is coming out soon for SAS! Visit:

methodology.psu.edu/ra/adap-treat-strat/qlearning

- These slides will be posted at

www-personal.umich.edu/~dalmiral/

Alter delivery mechanism, or improve motivation or form of treatment

Adherence is not a statistical nuisance; adherence indicates need to tailor treatment.