SMART-Adaptive Treatment Strategies

Sanjay Shete

sshete@mdanderson.org

Adaptive Treatment Strategies-WHY

- Adaptive treatment strategies: new paradigm for treatment and long term management of chronic, relapsing disorders such as smoking, cocaine abuse, depression etc.
- Adaptive treatment strategies: the treatment level and type is repeatedly adjusted according to ongoing response.

Adaptive Treatment Strategies-WHY

- Response is optimized-treatment type and dosage is modified as a function of response to past treatment
- Major challenge: delayed effects
- Each individual may be randomized multiple times over time
- The goals of such trials is to develop adaptive treatment strategies and are ***not confirmatory***

Adaptive Treatment Strategies-example

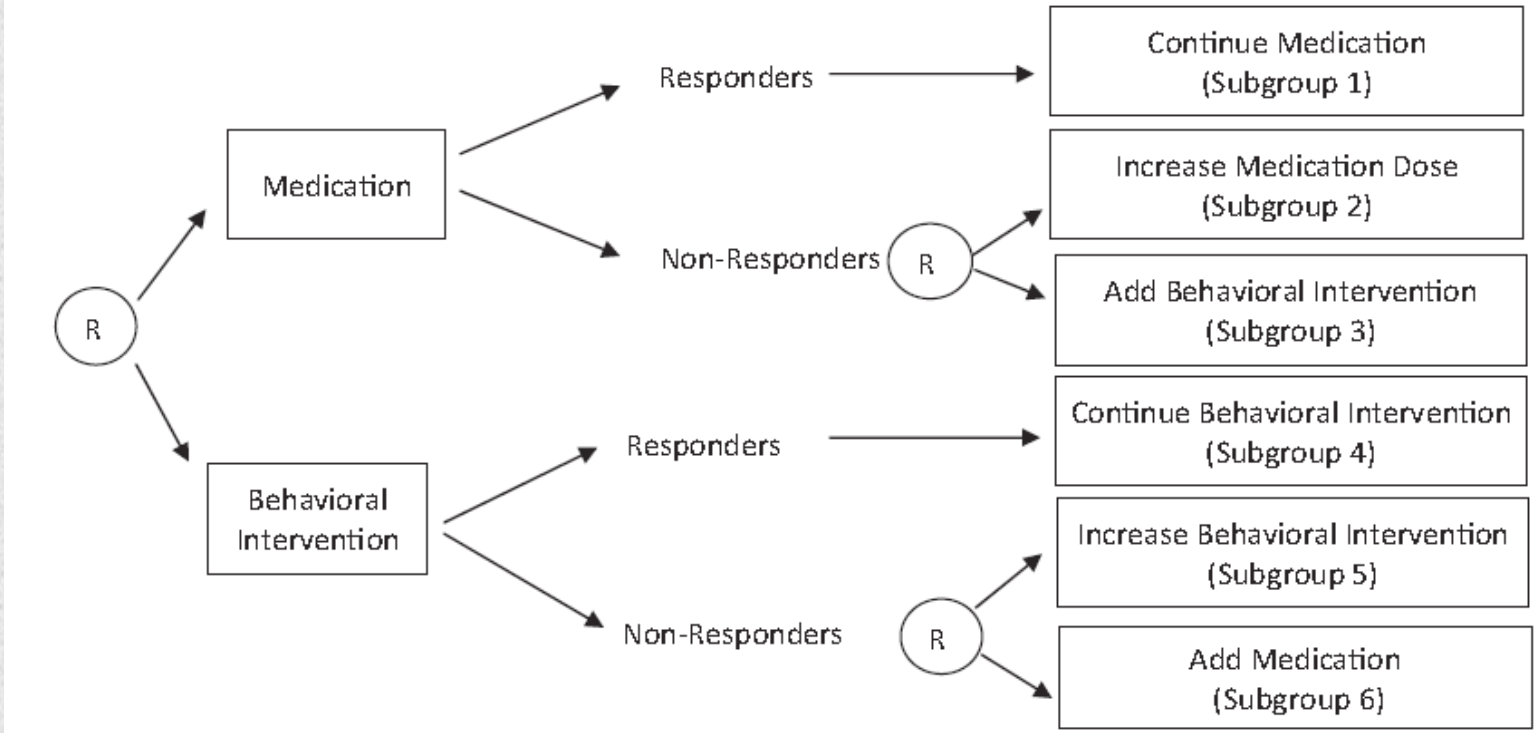

Adaptive Treatment Strategies-WHY

- The first decision rule can use the pre-treatment conditions such as level of addiction, age, years smoked etc.
- Intermediate response: reduction in number of cigarettes smoked per day or abstinence
- The second stage decision rule can use other intermediate outcomes such as adherence to initial treatment, self-management skills, motivation, depression etc.

Decision Rules and Tailoring Variables

- Decision rules are set to guide practitioners in deciding which intervention options use at each stage of the adaptive intervention, using available information relating to characteristics and ongoing performance of the participants.
- Tailoring variables: moderate effects of intervention. Types and dosage of intervention should be tailored according values of these variables.
- Not all moderators are tailoring variables: both sexes respond to high intensity behavioral intervention, however, it is found to be more beneficial to women. In this case, even though the sex is moderator both men and women should be offered the intervention because both groups are likely to benefit.

Adaptive Intervention for ADHD Children-table 1

Four Adaptive Interventions and Decision Rules Based on the ADHD Example

Adaptive intervention	Decision rule
(1, −1) First, offer low-intensity behavioral intervention; then add medication for nonresponders and continue low-intensity behavioral intervention for responders.	First-stage intervention option = {BMOD} IF evaluation = {nonresponse} THEN second-stage intervention option = {AUGMENT} ELSE continue on first-stage intervention option
(−1, −1) First, offer low-dose medication; then add behavioral intervention for nonresponders and continue low-dose medication for responders.	First-stage intervention option = {MED} IF evaluation = {nonresponse} THEN second-stage intervention option = {AUGMENT} ELSE continue on first-stage intervention option
(1, 1) First, offer low-intensity behavioral intervention; then increase the intensity of behavioral intervention for nonresponders and continue low-intensity behavioral intervention for responders.	First-stage intervention option = {BMOD} IF evaluation = {nonresponse} THEN second-stage intervention option = {INTENSIFY} ELSE continue on first-stage intervention option
(−1, 1) First, offer low-dose medication; then increase the dose of medication for nonresponders and continue low-dose of medication for responders.	First-stage intervention option = {MED} IF evaluation = {nonresponse} THEN second-stage intervention option = {INTENSIFY} ELSE continue on first-stage intervention option

Adaptive Intervention for ADHD Children

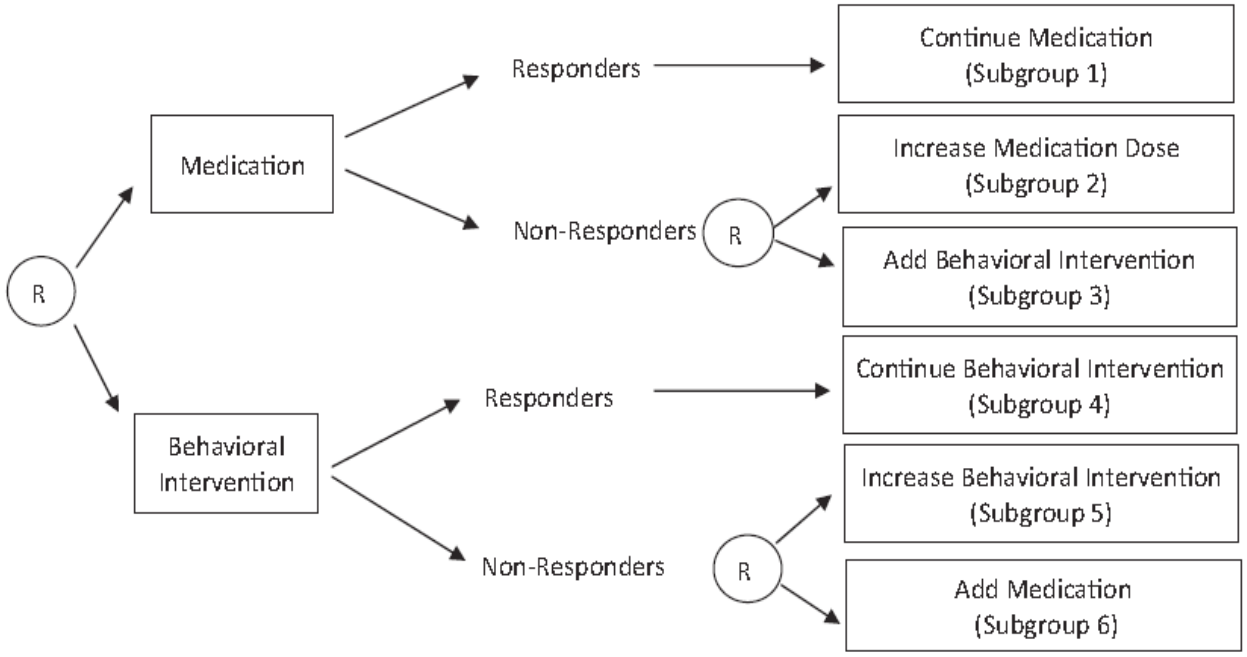

Figure 1. Sequential multiple assignment randomized trial for attention-deficit/hyperactivity disorder (ADHD) study.

Research Questions of Interest

1. The comparison of different intervention options at different stages of intervention (e.g. the difference between the first-stage intervention options or difference between second-stage intervention options for non-responding participants.
2. The comparison of adaptive interventions (i.e. sequence of decision rules) that are embedded within the SMART (e.g. 4 adaptive interventions referred in slide 7).

ADHD study

- In the first stage children were randomly (p=0.5) assigned to low dose medication or a low dose of behavioral intervention. Beginning at 8 weeks, response to first stage intervention was evaluated monthly until the end of school year.
- At each monthly assessment, children whose average performance on IBT and who were rated by teachers as impaired on the IRS were considered as non-responders to the first stage intervention and were randomly (0.5) assigned to one of the second stage intervention (i.e. increase the medication dose or augment the first stage intervention with other type of intervention.

Adaptive Interventions

- Each adaptive intervention embedded in this design is operationalized by a sequence of decision rules that specifies the intervention options at each stage for both responders and non-responders. Accordingly, there are responding and non-responding children who are ***consistent*** with each adaptive intervention.
- Children in subgroups 1 and 2 are consistent with an adaptive intervention, (-1, 1), begin with a low dose of medication and then non-responders receive an increased dose of medication (subgroup 2), whereas responders continue with the same low dose medication (subgroup 1).

Types of SMART Designs

- SMART designs vary in the extent and form of the tailoring that is incorporated in the design.
- This has to be based on scientific evidence as well as ethical considerations.

SMART-without embedded tailoring variables

- Participants are first randomized to two different first-stage intervention options (B or C).
- After some time (e.g. six months, 12 months etc.) all participants are re-randomized to two second-stage options (D or E) regardless of any *intermediate outcome or prior treatment*.
- Because there are no embedded tailoring variables in the SMART design, the interventions are non-adaptive.

SMART-without embedded tailoring variables

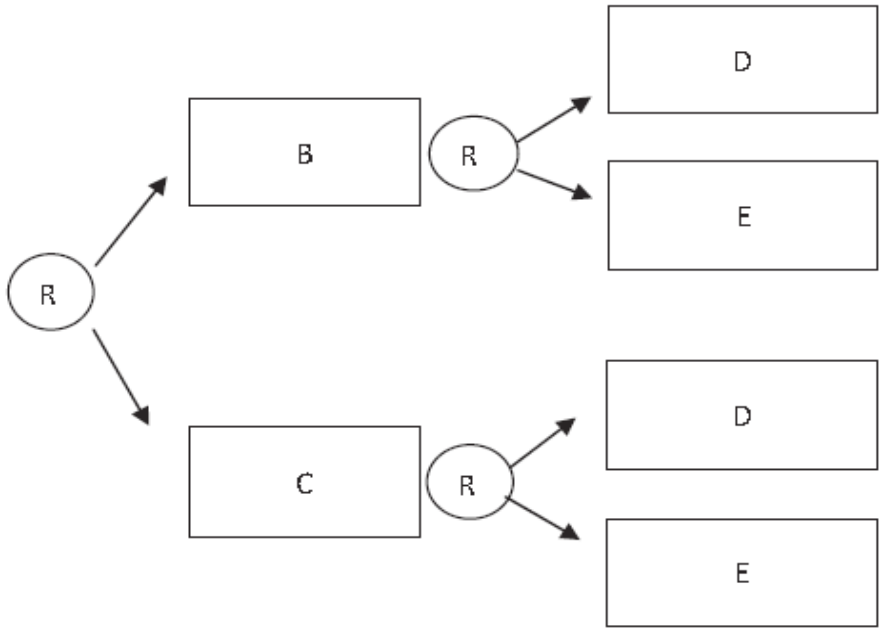

Figure 2. Sequential multiple assignment randomized trial (SMART) with no embedded tailoring variables.

SMART-everyone is re-randomized

- First, individuals are randomized to first stage interventions B or C.
- After some time (e.g. 12 months), all participants are re-randomized to second stage intervention options that depends on the intermediate outcome.
- Responders are randomized into maintenance interventions (M or M+)
- Non-responders are randomized into a third intervention (E) or receive the combined intervention (B+C)
- There are 8 adaptive interventions embedded in this design (1) begin with B and then offer E to non-responders and M for responders (2) begin with B and then offer E to non-responders and M+ for responders and so on.

SMART-everyone is re-randomized

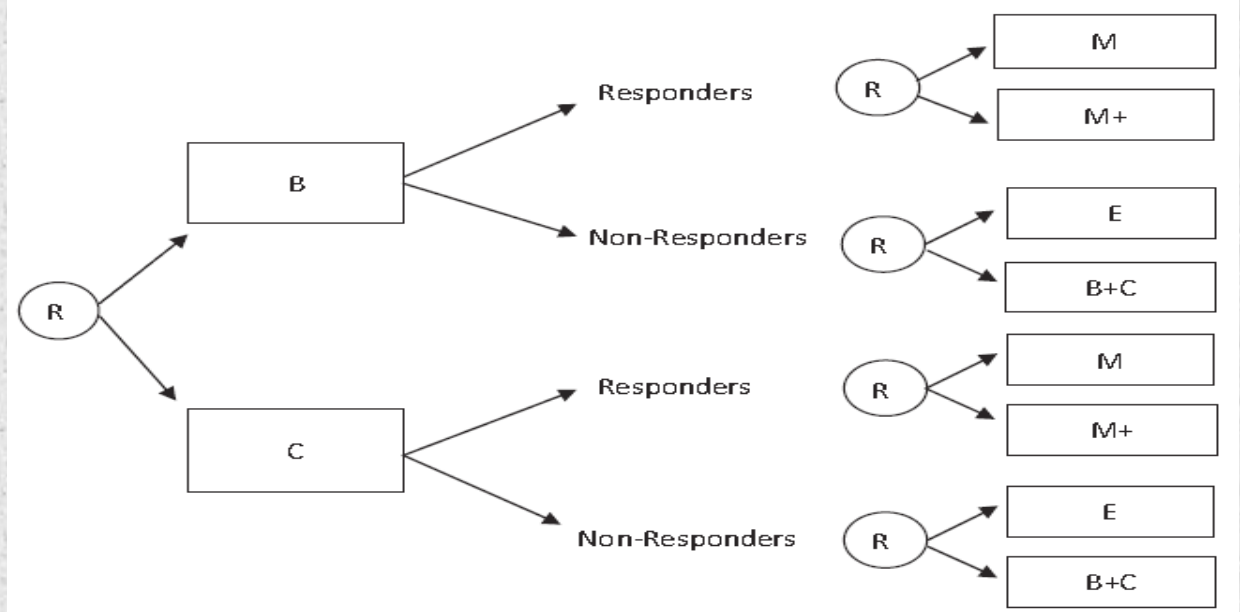

Figure 3. Sequential multiple assignment randomized trial (SMART) design in which participants are rerandomized to different second-stage intervention options depending on an intermediate outcome (e.g., response/nonresponse).

SMART-re-randomization depends on intermediate outcome and prior treatment

- First, individuals are randomized to first stage interventions B or C.
- Only non-responders to B are re-randomized to B+ or B+C.
- Responders to B (or C) are not re-randomized and stay on the B (or C).
- Non-responders to C are not re-randomized and are offered C+, an intensified version of C.
- There are 3 adaptive interventions embedded in this design (1) begin with B and then offer B+ to non-responders and responders stay on B (2) begin with B and then offer B+C to non-responders and responders stay on B, (3) begin with C and responders stay on C and non-responders receive C+.

SMART-re-randomization depends on intermediate outcome and prior treatment

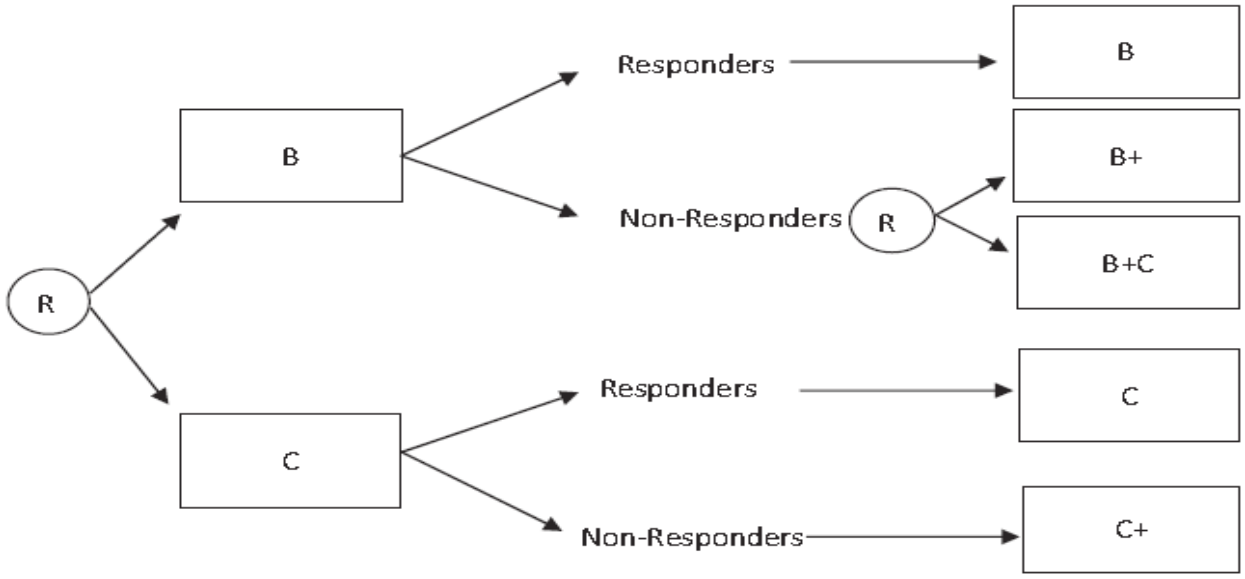

Figure 4. Sequential multiple assignment randomized trial (SMART) design in which the decision whether to rerandomize depends on an intermediate outcome and prior treatment.

Analyses Methods

- Observational data (O1, A1, O2, A2, Y)

1. O1 and O2 are vectors of pre-treatment information and intermediate outcomes, respectively (e.g. O1: severity of ADHD at the baseline, number of cigs/day, sex, age etc. O2: response status, reduction in number of cigs smoked/day, adherence to first stage intervention
2. A1 and A2 are first- and second-stage intervention options (e.g. A1 indicator: 1=low-intensity Behavioral intervention and -1=low-dose medication, A2NR: 1=increase the initial intervention; -1=augment the initial intervention with the other type of intervention)
3. Y is outcome of interest (e.g. smoking cessation, reduction in number of cigs smoked/day, evaluation of child's performance at the end of school year)

Questions of Interests

- The multiple randomizations in the SMART allow the investigator to estimate a large variety of causal effects important in the development of adaptive interventions.
- Typical focus is (a) difference between first stage intervention options, (b) difference between second stage intervention options, and (c) comparison of adaptive interventions that are imbedded within the SMART design.

Comparing first-stage intervention options

- In the context of specified second stage intervention options, does starting with low-intensity behavioral intervention result in a better long-term outcome relative to starting with low-dose medication? This question is addressed by pooling Y from subgroups 1-3 and comparing resulting average to the pooled Y from subgroups 4-6.
- SMART design investigators call it main effects of first stage intervention. This seems to be very problematic!

Comparing second-stage intervention options

- Among those who do not respond to their initial intervention, is there a difference between intensifying initial intervention versus augmenting the initial intervention?
- Figure 1, pool Y from subgroups 2 and 4 and compared the average to pooled Y from 3 and 6.
- SMART design investigators call it main effects of second stage intervention option for non-responding children. This seems to be very problematic!

Comparing adaptive interventions embedded within SMART

- Consider estimating mean outcome of only one of the four adaptive interventions (1,1).
- (1,1): take average of Y in subgroups 4 and 5. **INCORRECT**!
- Note that the outcomes of all children responding to behavioral intervention are included in the sample average, but the outcomes of only non-responders to behavioral intervention are included in this sample average.
- The sample average of subgroups 4 and 5 is biased estimator for the mean outcome of adaptive intervention (1,1).
- The bias occurs because, non-responding children are re-randomized and thus split into two subgroups (5 and 6), whereas the responding children are not re- randomized and thus not split into two subgroups.
- Sample average contains over-representation of outcomes from responding children and under-representation of outcomes from non-responding children.

Inverse-probability-of-treatment weights

- To accommodate the over/under representation, weights can be used. One choice of weights that is commonly used in the estimation of marginal structural models is inverse-probability-of-treatment weights.
- In our case W= 2 for responders and W=4 for non-responders. That is, each participant receives a weight that is inversely proportional to his/her probability of receiving his/her own adaptive intervention.
- This is simply taking weighted average instead of simple average.

Weighted Average

- Average of weighted outcomes of all children who are consistent with adaptive intervention (1,1): $\frac{\sum_{i=1}^{N(1,1)} w_i Y_i}{\sum_{i=1}^{N(1,1)} w_i}$
- w_i is the weight assigned to each individual who is consistent with adaptive intervention (1,1), N(1,1) is the number of children consistent with adaptive intervention (1,1).
- One may standardize weights so that the sum is equal to 1.

Weighted Average

- Note: the weight assigned to each child depended on whether the child was a responder or not to the first-stage intervention.
- That means the distribution of weights depends on the observed response rate in the sample, a statistic that varies from one sample to another.
- In order to account for sample to sample variance in the distribution of the weights, robust (sandwich) standard errors can be used to make inference (e.g. obtain p-value, CI etc.)

Analyses Plan

- One can obtain average weighted average, separately, for each of the four adaptive interventions that are embedded within the SMART. Then, test for the equality of means.
- It is better to do a single analyses in regression framework. One can adjust for the baseline factors that may be correlated with outcome of interest…accurate estimation of effects associated with adaptive interventions.

Analyses-Regression Framework

$$E[Y|A_1, A_{2NR}, O_1] = \beta_0 + \beta_1 A_1 + \beta_2 A_{2NR} + \beta_3 A_1 A_{2NR} + \gamma^T O_1,$$

- Y is the school performance at the end of school year, O1 is baseline covariates (measured prior to first-stage intervention), β0 is intercept, β1, β2 and β3 are regression coefficient for first-stage intervention, the second-stage intervention offered to non-responders, and the interaction between them; and γis the regression coefficient for the baseline covariates O1.
- E(Y|(1,-1)) = β0+ β1- β2- β3 is the average school performance of children on the adaptive intervention (1,-1).

Analyses-Regression Framework

- $E(Y|(-1,-1)) = \beta 0 - \beta 1 - \beta 2 + \beta 3$ is the average school performance of children on the adaptive intervention (-1,-1).
- $E(Y|(1,1)) = \beta 0 + \beta 1 + \beta 2 + \beta 3$ is the average school performance of children on the adaptive intervention (1,1).
- $E(Y|(-1,1)) = \beta 0 - \beta 1 + \beta 2 - \beta 3$ is the average school performance of children on the adaptive intervention (-1,1).

Analyses-Regression Framework

- Aim of this regression analyses is to compare the four adaptive interventions embedded in the SMART design, therefore, it does not include interactions between A1, A2NR, and baseline factors or intermediate outcome (O1, O2).
- Although, one could include interacts of O1 and intervention options to explore whether these interventions should be further tailored.
- However, inclusion of O2 or its interaction with interventions can lead to bias (why!)

Analyses-Regression Framework

- The parameters in equation can be

$$E[Y|A_1, A_{2NR}, O_1] = \beta_0 + \beta_1 A_1 + \beta_2 A_{2NR} + \beta_3 A_1 A_{2NR} + \gamma^T O_1,$$

estimated using the SAS GENMOD procedure by minimizing:

$$\sum_{j=1}^{M} W_j (Y - \beta_0 - \beta_1 A_{1j} - \beta_2 A_{2j} - \beta_3 A_{1j} A_{2NRj} - \gamma^T O_{1j})^2$$

to estimate regression coefficients. Standard errors are estimated using the robust (sandwich) estimates provided by the SAS GENMODE procedure.

Analyses-SMART with no embedded tailoring variables

- A1 = -1 for B and +1 for C
- A2= -1 for D and +1 for E
- Weights are not required in this setting because everyone is randomized equal number of times (twice).

$$E[Y|A_1, A_2, O_1] = \beta_0 + \beta_1 A_1 + \beta_2 A_2 + \beta_3 A_1 A_2 + \boldsymbol{\gamma}^T O_1$$

SMART-everyone is re-randomized

- A1 = -1 for B and +1 for C
- A2R= -1 for M and +1 for M+
- A2NR=-1 for E and +1 for B+C
- Weights are not required in this setting because everyone is randomized equal number of times (twice).

$$E[Y|A_1, A_{2R}, A_{2NR}, O_1] = \beta_0 + \beta_1 A_1 + \beta_2 A_{2R} + \beta_3 A_{2NR} + \beta_4 A_1 A_{2R} + \beta_5 A_1 A_{2NR} + \beta_6 A_{2R} A_{2NR} + \beta_7 A_1 A_{2R} A_{2NR} + \gamma^T O_1,$$

SMART-re-randomization depends on intermediate outcome and prior treatment

- A1 = -1 for B and +1 for C
- A2NRB= -1 for B+ and +1 for B+C
- W=4 for non-responders to B because these participants were randomized twice (each with equal probability of 0.5), W=2 for all other participants.
- A2NRB is nested within A1=-1. To represent this nested structure, consider an indicator variable Z that is equal to 0 when A1=1 (the initial intervention option is C) and equals to 1 when A1=-1 (the initial intervention option is B).

$$E[Y|A_1, A_2Z, O_1] = \beta_0 + \beta_1 A_1 + \beta_2 A_{2NRB} Z + \gamma^T O_1$$

ADHD data analysis

- 149 children (75% boys) between age 5-12 (mean age 8.6 years) participated in the study. After accounting for missing/dropout, the effective sample size was 139.
- At the first stage, 71 children were randomized to receive low-dose medication, and 68 were randomized to receive low-dose behavioral intervention.

ADHD data analysis

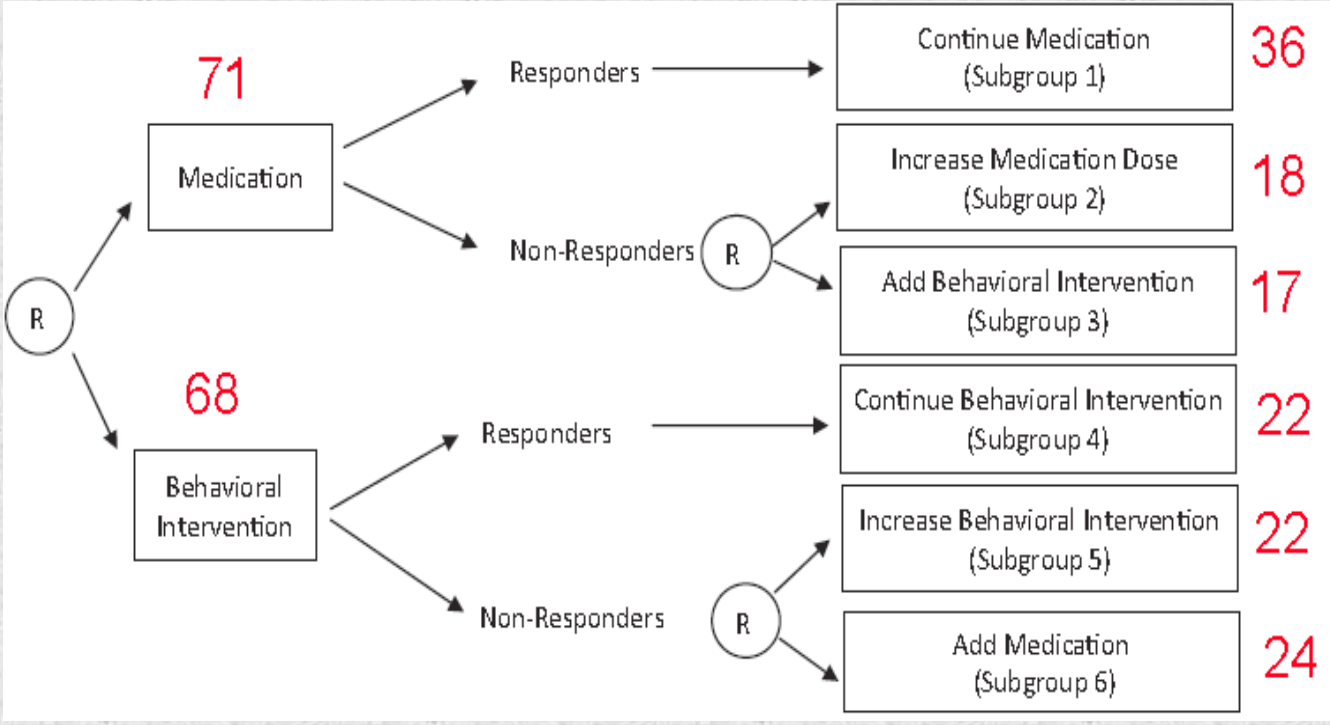

ADHD data analysis

- Outcome Y: Impairment Rating Scale after an 8-month period: range is 1 to 5, higher values >>better performance.
- Baseline Factors: Medication prior to the first-stage intervention, ADHD symptoms at the end of the previous school year: range 0 to 3, larger values >> fewer symptoms (i.e. better school performance).
- Oppositional defiant disorder (ODD): binary outcome.

ADHD data analysis- parameter estimates

Parameter	Estimate	Robust *SE*	95% confidence limit LL	95% confidence limit UL	Z	Pr > |Z|
Intercept	3.43	0.23	2.97	3.89	14.63	<.0001
Baseline: ODD diagnosis	0.37	0.18	0.02	0.72	2.07	.0384
Baseline: ADHD symptoms	0.57	0.14	0.29	0.85	3.95	<.0001
Baseline: Medication before Stage 1	−0.61	0.25	−1.10	−0.13	−2.47	.0134
A1	0.07	0.09	−0.11	0.24	0.75	.4555
A2	0.02	0.08	−0.13	0.18	0.26	.7924
A1 * A2	−0.12	0.08	−0.27	0.04	−1.46	.1436

ADHD data-mean for adaptive interventions

Estimated Mean and Standard Error for Each Adaptive Intervention

Adaptive intervention			Responders		Nonresponders		Estimated weighted mean	Robust *SE*
	Stage 1	Stage 2	Sample size	Sample mean	Sample size	Sample mean		
(1, −1)	BMOD	AUGMENT	22	4.64	24	4.08	4.36	0.15
(−1, −1)	MED	AUGMENT	36	4.39	17	3.47	4.00	0.15
(1, 1)	BMOD	INTENSIFY	22	4.64	22	3.96	4.17	0.22
(−1, 1)	MED	INTENSIFY	36	4.39	18	4.22	4.27	0.13

ADHD data analysis

Estimated Differences Between the Four Adaptive Interventions Based on the Estimated Regression Coefficients in Table 2

Label	Estimate	95% confidence limit		Robust *SE*	χ^2	Significance
		LL	UL			
Difference between (1, −1) and (−1, −1)	0.36	−0.06	0.79	0.22	2.82	.0932
Difference between (1, 1) and (−1, 1)	−0.10	−0.60	0.40	0.26	0.16	.6920
Difference between (1, −1) and (1, 1)	0.19	−0.32	0.70	0.26	0.55	.4600
Difference between (1, −1) and (−1, 1)	0.09	−0.31	0.49	0.20	0.20	.6563
Difference between (−1, −1) and (1, 1)	−0.17	−0.70	0.35	0.27	0.42	.5161
Difference between (−1, −1) and (−1, 1)	0.27	−0.08	0.63	0.18	2.26	.1328

Results-ADHD

- None of the adaptive interventions are better than other interventions.
- Significance should be tested for the highest mean versus the next highest mean.
- First-stage, second-stage interventions and their interactions were not significant.
- Sample size is relatively small.
- Power of the study may be very limited. (N=169)
- http://methodologymedia.psu.edu/smart/samplesize

- *Reference: Nahum-Shani I, Qian M, Almirall D, Pelham W, Gnagy B, Fabiano G, Waxmonsky J, Yu J, Murphy S. Experimental Design and Primary Data Analysis Methods for Comparing Adaptive Interventions. Psychological Methods, 2012 4:457-477.*

Q-Learning

- An approach to construct decision rules that operationalize optimal adaptive interventions.
- By using the Q-Learning technique to construct an adaptive intervention, we can find the sequence of decision rules that link the observed information concerning an individual to the most efficient intervention type and intensity or dosage.

Q-Learning

- Begin by finding the optimal decision rule at the second stage

  $d2^*(O1,a1,O2) = \arg\max Q2(O1,a1,O2,a2)$,

  where $Q2(O1,a1,O2,a2)=E[Y|O1,a1,O2,a2]$

- Q2 is the conditional expectation that provides the expected outcome of choosing second stage option a2, given the information available (O1,a1,O2)

Q-Learning

- Move backwards in time to construct the optimal decision rule at the first stage

  $d1^*(O1)=\arg\max Q1(O1,a1)$

  where

$Q1(O1,a1)=E[\max_{a2} Q2(O1, a1,O2,a2)|O1,a1]$

Q-Learning

- Assumption:
- linear models for Q1 and Q2 are correct and the observations from one individual to another are independent.
- the estimators of the regression coefficients are consistent(unbiased in large samples) for the true regression coefficients.
- sample size is sufficiently large so that the distribution of the estimators for the regression coefficients can be well approximated by the normal distribution.

Q-Learning

- Q-Learning can be used to estimate the optimal sequence of decision rules in a straightforward and intuitive manner.
- It appropriately controls for the optimal 2nd-stage intervention option when assessing the effect of the 1st-stage intervention.
- The effects estimated by Q-Learning incorporate both the direct and indirect effects of the 1st-stage intervention options, the combination of which is necessary for making intervention decision rules.
- Q-Learning reduces potential bias resulting from unmeasured causes of both the tailoring variables and the primary outcome.
- The article illustrated the application of Q-Learning using a simplified version of the Adaptive Interventions for children with ADHD study, with the general aim to guide researchers who wish to apply this method to construct high-quality adaptive interventions

V-SMART

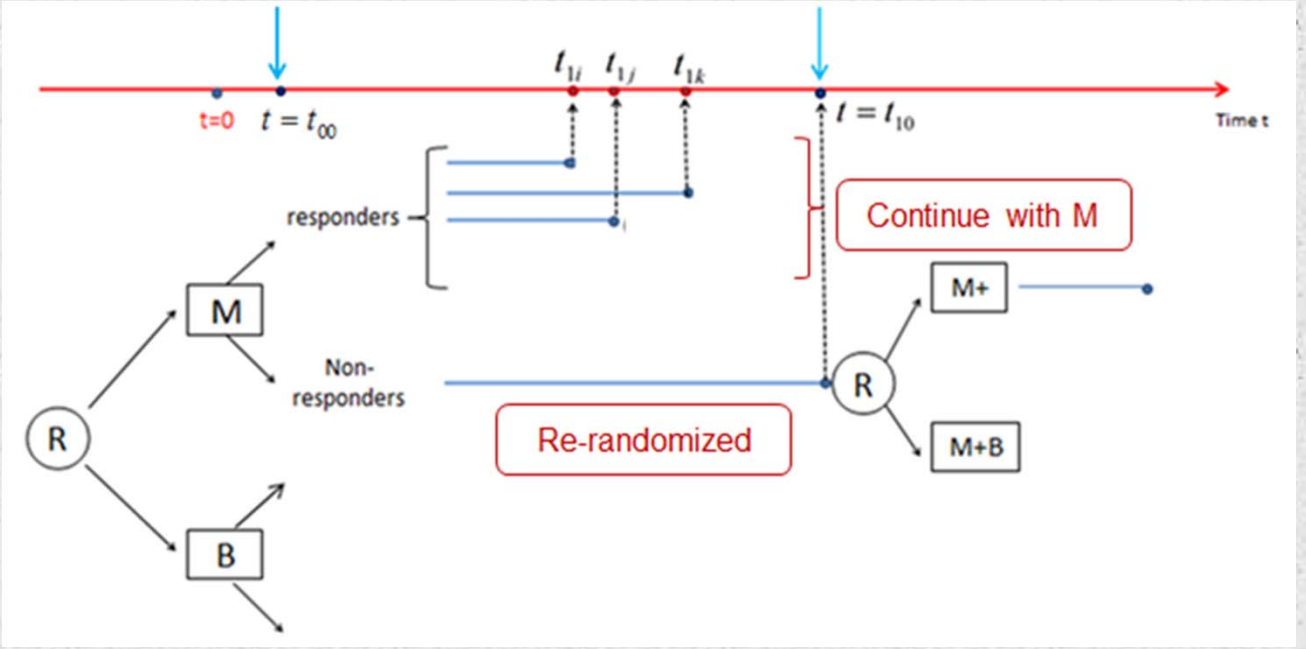
t=0
$t = t_{00}$
t_{1i}
t_{1j}
t_{1k}
$t = t_{10}$
Time t
responders
Continue with M
M
M+
Non-
responders
R
R
Re-randomized
M+B
B

V-SMART

- Two stage time varying SMART design
- Goal is to avoid side effects
- Treatment is given till time t00. Response is continuously monitored
- Responders are re-randomized as soon as they reach pre-specified threshold (e.g. number of cigs/day)
- Responders are immediately put on maintenance regime as opposed to continuing initial regimen till a fixed time point t10

SMART

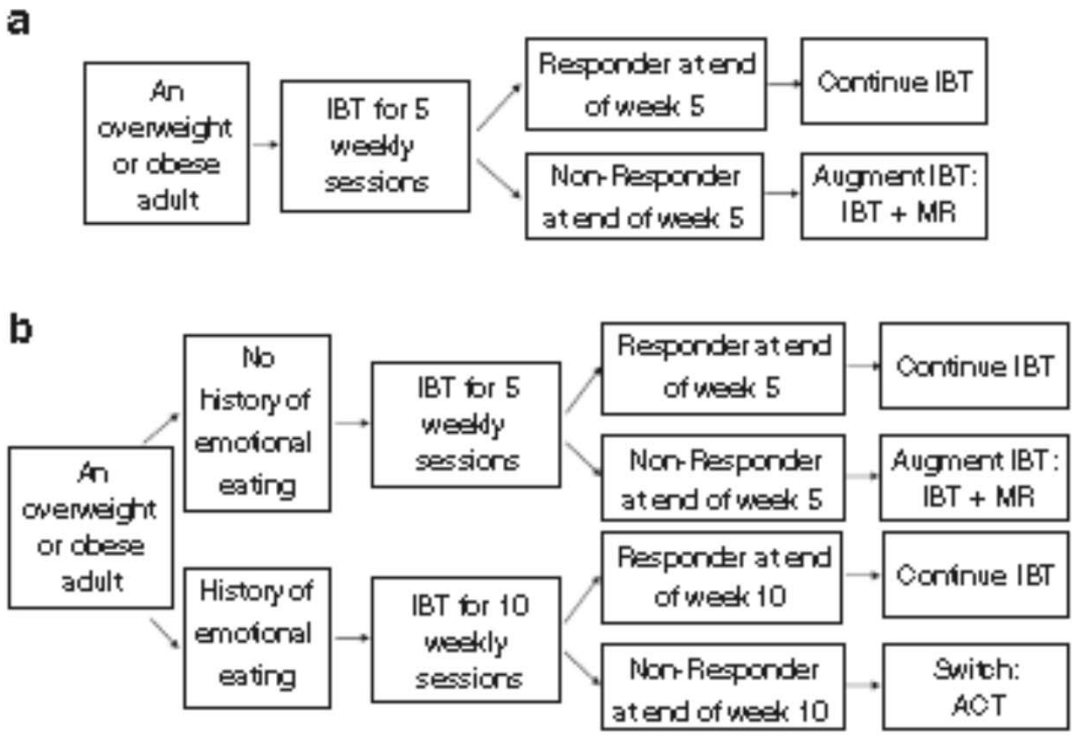
a
An overweight or obese adult
IBT for 5 weekly sessions
Responder at end of week 5
Continue IBT
Non-Responder at end of week 5
Augment IBT: IBT + MR
b
An overweight or obese adult
No history of emotional eating
IBT for 5 weekly sessions
Responder at end of week 5
Continue IBT
Non-Responder at end of week 5
Augment IBT: IBT + MR
History of emotional eating
IBT for 10 weekly sessions
Responder at end of week 10
Continue IBT
Non-Responder at end of week 10
Switch: ACT