SAS PROC QLEARN Users' Guide Version 1.0.0

Ashkan Ertefaie
Daniel Almirall
Liying Huang
John J. Dziak
Aaron Wagner
Susan Murphy

Please send questions and comments to *MChelpdesk@psu.edu*.

The development of PROC QLEARN was supported by the National Institute on Drug Abuse Grant P50-DA10075 to The Center for Prevention and Treatment Methodology. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute on Drug Abuse or the National Institutes of Health.

The suggested citation for this users' guide is

Ertefaie, A., Almirall, D., Huang, L., Dziak, J. J., Wagner, A. T., & Murphy, S. A. (2012). *SAS PROC QLEARN users' guide* (Version 1.0.0). University Park: The Methodology Center, Penn State. Available from http://methodology.psu.edu.

Contents

1 Overview of PROC QLEARN 3
1.1 Macro Features 3
2 Introduction to Adaptive Interventions and SMART 4
2.1 Adaptive Interventions 4
2.2 Sequential, Multiple Assignment, Randomized Trial (SMART) 5
2.3 Data Structure 5
2.4 Example of a SMART Design 6
2.5 Q-learning Overview 7
3 Q-learning: Statistical Models 8
3.1 Regression Models 9
3.2 Contrast Matrix for First-Stage Regression 10
4 Technical Details 13
4.1 Estimation procedures 13
4.2 Missing Data 13
4.3 Confidence Intervals for First-Stage Parameters and Contrasts 13
5 Running PROC QLEARN 15
5.1 PROC QLEARN Syntax 15
5.2 Invoking the QLEARN Procedure 15
5.3 Options for Input 15
5.4 Required Statements for PROC QLEARN 16
5.5 Optional Statements for PROC QLEARN 17
References 20
Appendices 23
Appendix A - Adaptive Interventions for Children with ADHD 23
Appendix B: Adaptive Interventions for Minimally Verbal Children 30

1 Overview of PROC QLEARN

PROC QLEARN is a SAS procedure that implements a generalization of Q-learning, a method developed in computer science, to inform the development of adaptive interventions. In an adaptive intervention, treatment is individualized based on personal characteristics. Using a sequence of decision rules, treatment is adjusted over time in response to an individual's ongoing performance, status, or personal circumstances.

Q-learning is an extension of the standard regression method that can be used with longitudinal data in which treatments vary over time. In this documentation, we describe how to use PROC QLEARN with data from a sequential, multiple assignment, randomized trial (SMART). SMARTs are clinical trial designs that generate high-quality data explicitly for the purpose of developing adaptive interventions.

1.1 Macro Features

PROC QLEARN is developed for SAS[1] for Windows. The PROC is designed for situations where the outcome is continuous, there are two decision stages and up to two treatment options at each decision stage. Specifically, PROC QLEARN allows users to use a two-decision-stage Q-learning algorithm to construct an adaptive intervention using data arising from a SMART study. This implementation of PROC QLEARN uses a bootstrap procedure to obtain confidence intervals for the estimated Q-learning regression parameters from decision stage 1.

[1]SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries.

2 Introduction to Adaptive Interventions and SMART

2.1 Adaptive Interventions

An adaptive intervention is a sequence of decision rules that specify whether, how, or when to alter the intensity, type, or delivery of treatment at decision stages in the medical care process. Adaptive interventions aim to take advantage of individual heterogeneity in response to treatments in order to maximize health outcomes. They do this by adapting and re-adapting treatments over time based on individuals characteristics and evolving health status. Adaptive interventions are also known as dynamic treatment regimes (Robins, 1986; Murphy et al., 2001)adaptive treatment strategies (Lavori et al., 2000; Murphy, 2005), multi-stage treatment strategies (Thall et al., 2002) and treatment policies (Lunceford et al., 2002; Wahed & Tsiatis, 2004; Wahed & Tsiatis, 2006).

There are three important parts at each decision stage of an adaptive intervention: treatment options at each stage, tailoring variables, and decision rules. Tailoring variables are variables that serve as inputs to the decision rules; they are variables that influence responsiveness to (or can be used to determine the need for) treatment type or dosage. Decision rules use data from tailoring variables to match an individual's characteristics and status with a treatment option.

There are two types of tailoring variables: baseline and time-varying. Baseline tailoring variables may include age, comorbidities, race, gender, previous history of treatment, measures of severity, and contextual characteristics representing risk or protective factors. Baseline tailoring variables are observed prior to, and used in making the first decision. In general, time-varying tailoring variables are observed at each decision stage. In this PROC, since we only consider two stage adaptive interventions, time-varying tailoring variables are observed up to and at the end of first-stage treatment. They may include adherence to first-stage treatment, side effects, or response to first-stage treatment.

For example, suppose that an investigator is interested in evaluating treatment A vs B on a group of patients. Assume that the prior scientific evidence suggests that male patients

benefit from treatment A, while female patients benefit from treatment B. Also assume that the investigator conjectures that among non-responders to the first-stage treatment those who develop side effects might benefit from switching the first-stage treatment, while patients who do not develop any side effects might benefit from intensifying the first-stage treatment. In this hypothetical example,

- gender is the baseline tailoring variable,
- treatments A and B are first-stage treatment options,
- the presence of side effects is the time-varying tailoring variable,
- intensifying and switching are second-stage treatment options.

PROC QLEARN is used with data from a sequential, multiple assignment, randomized trial (SMART) to construct adaptive interventions. In the next section, we discuss the SMART.

2.2 Sequential, Multiple Assignment, Randomized Trial (SMART)

The sequential, multiple assignment, randomized trial (SMART) is a clinical trial design in which each subject can proceed through stages of treatment. At each decision stage, participants are randomized to one of the available treatment options at that stage. For example, at first-stage all participants may be randomized to treatment A vs treatment B. In the second stage (following the first-stage treatment) only (non-)responders to the first-stage treatment may be randomized. Randomizations may be restricted based on ethical concerns, feasibility, availability, or suitability of different sets of treatment options. In the next subsection, we present the typical data structure of the SMART design.

For more information on SMART, please see Murphy, (2005); Nahum-Shani, et al., (2010a); and Nahum-Shani, et al., (2010b).

2.3 Data Structure

The PROC accommodates longitudinal data arising from SMART with **two** decision stages. The observed data on each individual is given by $(O_1, A_1, O_2, S, A_2, Y)$. O_i, for $i = 1, 2$, is a set of covariates in the beginning of the i th stage. A_i denotes the treatment options in the i th

stage. The PROC requires two treatment options at each decision stage. S is a binary variable which is coded as one if an individual has been rerandomized at stage 2 and coded zero otherwise. Y is the primary outcome and assumed to be observed for everyone at the end of the study. This PROC requires continuous Y. The vectors O_1 and O_2 are called *baseline covariates* and *intermediate outcomes*, respectively. O_1 and O_2 may include candidate baseline and time-varying tailoring variables, respectively.

Note: Although contrast coding of A_i is not required, to facilitate interpretation, we use contrast coding (e.g. A_i is coded by 1 or -1) in this user guide.

The following is an example SMART design with two treatment options at each stage.

2.4 Example of a SMART Design

The example SMART below is modeled after a SMART designed to build an adaptive intervention to optimize school performance of children with attention deficit hyperactivity disorder (ADHD)(Pelham et al., 2002; Nahum-Shani et al., 2010a; and Nahum-Shani et al., 2010b). The data used here is simulated.

The data includes baseline covariates that are candidate baseline tailoring variables, including age, race and school performance in the previous year (O_1). In the first-stage, participants had a probability $1/2$ of assignment to either medication or behavioral modification treatment (A_1). After eight weeks, children's responses to the first-stage treatment was evaluated monthly until the end of the school year. The monthly assessment is an intermediate outcome (O_2) which can be considered as a candidate time-varying tailoring variable. Based on the information obtained by the monthly assessments, children were categorized as responders or non-responders to the first-stage treatment. In the second-stage, responders to the first-stage treatment continued receiving the same treatment, while non-responders had a probability $1/2$ of being randomized to either an intensified version of the first-stage treatment or augmentation of the first-stage treatment with the other treatment option (A_2) (i.e. $S=1$ for non-responders and zero otherwise). The school performance at the end of the school year (Y), coded so that high values are preferable, is the outcome.

By design, there are four adaptive interventions embedded in this SMART design.

- Start with medication. If the patient is non-responsive, increase intensity of medication; if the patient is responsive, continue with the same medication dose.
- Start with medication. If the patient is non-responsive, assign to behavioral modification; if the patient is responsive, continue with the same medication dose.
- Start with the behavioral modification treatment. If the patient is non-responsive, assign intensified behavioral modification; if the patient is responsive, continue the same behavioral modification.
- Start with the behavioral modification. If the patient is non-responsive, assign medication; if the patient is responsive, continue the same behavioral modification.

2.5 Q-learning Overview

Q-learning (Watkins, 1989; Watkins & Dayan, 1992) is a method generalized from computer science to extract the information obtained by a SMART design to form a sequence of individualized treatments that results in a desired outcome. In other words, Q-learning can be used with data from a SMART to help build an optimal adaptive intervention.

Q-learning starts from the second-stage and searches for a treatment option that optimizes the final outcome. Then given the optimally-chosen second-stage treatment, Q-learning moves backward to the first-stage and searches for a treatment option that optimizes the outcome (Murphy, 2003). The backward induction used by Q-learning avoids treatment options that appear to be optimal in the short term but may lead to an undesirable or less desirable final outcome. In the following section, we explain how Q-learning can be used to construct an optimal sequence of decision rules.

3 Q-learning: Statistical Models

Q-learning fits a linear regression model at each decision stage. The regressions are fit sequentially starting from the second-stage, working backward to the first-stage. At each stage, the purpose of the regression is to evaluate usefulness of candidate tailoring variables and to find the treatment option at that stage that maximizes the outcome. We explain the Q-learning procedure on studies with two stages because PROC QLEARN accommodates these designs.

To fit the second-stage model, only data from participants who are re-randomized is used. The primary outcome is regressed on the intermediate outcomes, assigned treatments at first- and second-stages and baseline covariates to fit the second-stage model. Using the estimated parameters in the second-stage model, the second-stage treatment option that optimizes the expected value of the primary outcome can be derived (see below). Next data from all participants is used to fit the first-stage model; this model is fit by regressing the appropriate outcome (see below) on the baseline covariates and the treatment assigned at the first-stage. Through the rest of this section, we present the regression models in more detail.

Q-learning uses the observed history of each participant at each stage, e.g. as $H_1 = O_1$ and $H_2 = (O_1, A_1, O_2)$ to construct the optimal decision rules d_1 and d_2 for the first- and second-stages respectively, as follows:

$$d_2(H_2) = \arg\max_{a_2} Q_2(H_2, a_2),$$
$$d_1(H_1) = \arg\max_{a_1} Q_1(H_1, a_1),$$

where

$$Q_2(H_2, a_2) = E[Y \mid H_2, a_2],$$
$$Q_1(H_1, a_1) = E[\max_{a_2} Q_2(H_2, a_2) \mid H_1, a_1],$$

(see Charkaborty, et al., 2010). In fact, $d_2(H_2)$ is the optimal treatment, which maximizes the expected value of the primary outcome given the covariate history H_2. $d_1(H_1)$ is the first-stage treatment that given covariates H_1, leads to an optimal outcome when combined with the optimal second-stage treatment, $d_2(H_2)$.

3.1 Regression Models

Regression models can be used to estimate the two functions Q_1 and Q_2. For simplicity we assume the following linear regression models:

1. The second-stage model (fit among individuals who are rerandomized at stage 2):

$$Q_2(H_2, A_2; \beta_{21}, \beta_{22}) = \beta_{21}^T H_{21} + \beta_{22}^T H_{22} A_2, \quad (1)$$

where H_{21} and H_{22} are vectors of covariates from the participant's history at the second-stage, H_2. H_{21} includes a "1" as the first element; this is the intercept for the model. H_{22} also includes "1" as the last element, so that the last parameter in β_{22} represents the main effect of the second-stage treatment. The parameters of this model are estimated by regressing the outcome Y on $(H_{21}, H_{22}A_2)$. This results in estimators, $\hat{\beta}_{21}, \hat{\beta}_{22}$. The estimated interaction parameters $\hat{\beta}_{22}$ are crucial to identifying the tailoring variables. Since, in this PROC, we are using contrast coding (-1,+1) for A_2, the estimated optimal treatment option given H_2 at stage 2 is as follows: if $\hat{\beta}_{22}H_{22} > 0$ then the treatment coded 1 is the best; whereas if $\beta_{22}H_{22} < 0$ then the treatment coded -1 is the best.

From this regression, the estimated outcome, Y, if the optimal treatment $d_2(H_2)$ were taken is,

$$\hat{Y}^{opt} = \max_{a_2} Q_2(H_2, a_2; \hat{\beta}_{21}, \hat{\beta}_{22}) = \hat{\beta}_{21}^T H_{21} + \max_{a_2}\{\hat{\beta}_{22}^T H_{22} a_2\}.$$

Since $A_2 \in \{-1, +1\}$, the preceding equation can be written as

$$\hat{Y}^{opt} = \max_{a_2} Q_2(H_2, a_2; \hat{\beta}_{21}, \hat{\beta}_{22}) = \hat{\beta}_{21}^T H_{21} + |\hat{\beta}_{22}^T H_{22}|. \quad (2)$$

The dependent variable for the first-stage regression, $\tilde{Y}$, is

$$S\hat{Y}^{opt} + (1-S)Y. \quad (3)$$

In other words, if a participant is randomized at stage 2, we set $\tilde{Y} = \hat{Y}^{opt}$ and if not, we set $\tilde{Y} = Y$.

2. The first-stage model (fit using data from all participants):

$$Q_1(H_1, A_1; \beta_{11}, \beta_{12}) = \beta_{11} H_{11} + \beta_{12} H_{12} A_1,$$

where H_{11} and H_{12} are vectors of covariates from the participant's baseline data, H_1. H_{11} includes a "1" as the first element; this is the intercept for the model. H_{12} also includes "1" as the last element, so that the last parameter in β_{12} represents the main effect of the first-stage treatment. The parameters of this model are estimated by regressing the outcome $\tilde{Y}$ on $(H_{11}, H_{12} A_1)$. This results in estimators, $\hat{\beta}_{11}$, $\hat{\beta}_{12}$. Similar to the second-stage, the estimated optimal treatment at the first-stage is: if $\hat{\beta}_{12} H_{12} > 0$ then the treatment is coded 1 is the best; whereas if $\hat{\beta}_{12} H_{12} < 0$ then the treatment coded -1 is the best.

The regression parameters in Q_j, for $j = 1, 2$, are consistently estimated (unbiased in large samples) if the fitted models are correctly specified. In the technical details section, we discuss the bootstrap based confidence intervals for the first-stage parameters provided by this PROC.

3.2 Contrast Matrix for First-Stage Regression

The PROC allows users to specify a contrast matrix. This matrix is used to estimate a "contrast" of the stage 1 regression parameters. A contrast is a linear combination of the stage 1 regression parameters. For example, a contrast can be used to represent the mean outcome for individuals with a particular covariate values or can be used to represent differences in the mean outcome between treatments for individuals with particular covariate values. The number of columns in a contrast matrix is equal to the number of regression parameters at the first-stage, and the number of rows is the number of contrasts that a user wishes to estimate.

As an example, suppose the first-stage estimated model is

$$\hat{\beta}_{11}H_{11} + \hat{\beta}_{12}H_{12}A_1,$$

where $H_{11} = (1, O_{11}, O_{12}, O_{13})$, $H_{12} = (O_{13}, 1)$ and A_1 is the treatment indicator at the first-stage (coded as -1 and +1). Since we have 6 parameters in this model, the contrast matrix must have 6 columns. In the following example we decide to estimate 4 contrasts so our contrast matrix has 4 rows. Let the 4×6 contrast matrix C_1 be

$$C_1 = \begin{pmatrix} 1 & 0 & 0 & 1 & -1 & -1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & -2 & -2 \\ 0 & 0 & 0 & 0 & 0 & -2 \end{pmatrix}$$

The first row of C_1 estimates $(1,0,0,1)\beta_{10} + (-1,-1)\beta_{12}$ which is the mean response, $\tilde{Y}$, among participants with $O_{11} = 0$, $O_{12} = 0$, $O_{13} = 1$, and $A_1 = -1$ (therefore, $A_1 O_{13} = -1$). The second row of C_1 estimates $(1,0,0,1)\beta_{11} + (1,1)\beta_{12}$ which is the mean response, $\tilde{Y}$, among participants with $O_{11} = 0$, $O_{12} = 0$, $O_{13} = 1$, and $A_1 = 1$ (therefore, $A_1 O_{13} = 1$). The third contrast estimates the difference between the first two contrasts. More specifically, the third row of the C_1 estimates

$[(1,0,0,1)\beta_{11} + (-1,-1)\beta_{12}] - [(1,0,0,1)\beta_{11} + (1,1)\beta_{12}] = (0,0,0,0)\beta_{11} + (-2,-2)\beta_{12}$ which is the difference in mean response between participants with $A_1 = -1$ and $A_1 = 1$ when $O_{13} = 1$. The last row estimates

$[(1,0,0,0)\beta_{11} + (0,-1)\beta_{12}] - [(1,0,0,0)\beta_{11} + (0,1)\beta_{12}] = (0,0,0,0)\beta_{11} + (0,-2)\beta_{12}$ which is the difference mean response between participants with $A_1 = -1$ and $A_1 = 1$ when $O_{13} = 0$.

Comparing the estimated contrast obtained from the first and the second rows of this matrix, allows us to estimate which first-stage treatment is best for individuals with $O_{11} = 0$, $O_{12} = 0$, $O_{13} = 1$. More specifically, if the mean response given $O_{11} = 0$, $O_{12} = 0$, $O_{13} = 1$, $A_1 = -1$ is greater than the mean response given $O_{11} = 0$, $O_{12} = 0$, $O_{13} = 1$, $A_1 = +1$, then the group of patients with characteristics $O_{11} = 0$, $O_{12} = 0$, $O_{13} = 1$ is estimated to benefit more from the treatment coded as $A_1 = -1$ than from the treatment coded as $A_1 = 1$.

The third and forth contrasts can be used to test if a group of patients benefits from one

treatment option as compared to the other treatment option. If the confidence interval corresponding to the contrast in the third row contains zero, then there is no evidence that patients with $O_{11} = 0$, $O_{12} = 0$, and $O_{13} = 1$ respond differently to the two different treatments. However, if the confidence interval lies entirely on the positive side of the real line, then we can conclude that the treatment coded as -1 leads to a better mean response among patients with $O_{11} = 0$, $O_{12} = 0$, and $O_{13} = 1$, than the treatment coded 1. The last row can be interpreted in a similar manner to the third contrast; however this contrast is for the group of patients with $O_{11} = 0$, $O_{12} = 0$, and $O_{13} = 0$.

The intermediate variable O_{13} can be used to tailor the first-stage treatment if the confidence intervals of the two estimated contrasts obtained from the third and last rows lie entirely on the different sides of the real line.

4 Technical Details

4.1 Estimation procedures

In PROC QLEARN, the regression coefficients of the fitted models are estimated using the least squares method, and the confidence intervals of the parameters are estimated using a bootstrap technique developed in Laber & Murphy (2011).

4.2 Missing Data

Missing covariates and/or missing primary outcome are permitted. Missing values should be represented as SAS system missing (".") as usual in SAS. The QLEARN procedure assumes that the data is missing completely at random (MCAR). The procedure eliminates any record with missing data on a covariate or outcome. That is, an individual with $S = 1$ will be ignored if there are any missing data in H_{21}, H_{22} or the primary outcome Y is missing. Any individual ($S = 1$ or $S = 0$) with missing data in H_{11}, H_{12} or missing the primary outcome Y will be ignored. Also, if the variable S is missing the individual will be ignored.

4.3 Confidence Intervals for First-Stage Parameters and Contrasts

Confidence intervals for the parameters in the first-stage model Q_1, see (1), requires a generalization of the bootstrap. Recall - see (2) and (3) - that the parameters in Q_1 are estimated by regressing $\tilde{Y}$ on covariates in H_1. Non-differentiability of the absolute value in (2) at zero (when $\beta_{22} H_{22} A_2 = 0$) can result in poor performance of the standard bootstrap-based confidence intervals (see Bickel & Freedman, 1981; Robins, 2004; Tsiatis (2006); and Chakraborty et al., 2010).

This PROC uses the *Adaptive Confidence Interval (ACI)* method developed by Laber et al., (2010) to construct the confidence intervals. The ACI is formed by partitioning the participants in sets for which the estimated stage 2 treatment effect (e.g., $\beta_{22} H_{22} A_2$) is near zero and for which the estimated stage 2 effect is far from zero and constructing smooth upper and lower bounds using the two partitions. The upper and lower bounds are calculated using a grid search method. The number of grid points and the limits of the grid search for each parameter are specified by

the statements *NGRID* and *GRID-SCALE,* respectively. The upper and lower bounds are then bootstrapped to form the confidence interval. As described in section 3.2, two sided hypothesis tests can also be performed for the parameters of the first-stage model by checking whether the corresponding confidence intervals contain zero.

5 Running PROC QLEARN

5.1 PROC QLEARN Syntax

The following statements are available in PROC QLEARN. The statements in bold are required; other statements are optional.

PROC QLEARN < *options for input* (see section 5.3 below)> ;
MAIN1 *variables* ;
TAILOR1 *variables* ;
MAIN2 *variables* ;
TAILOR2 *variables* ;
RESPONSE *variable* ;
STG1TRT *variable* ;
STG2TRT *variable* ;
STG2SAMPLE *variable* ;
NBOOTSTRAP1 *value* ;
NGRID *value* ;
GRID-SCALE *value* ;
ALPHA *value* ;
RUN;

5.2 Invoking the QLEARN Procedure

To invoke the QLEARN procedure, first use the following line of SAS code:

PROC QLEARN DATA = *SAS-data-set* < *options for input*> ;

The data file may contain more variables than those used in the analysis. It must contain 2 variables to be used as indicators for the first- and second-stage treatment options and at least one continuous variable to be used as an outcome (final outcome).

There are several options that may be specified in the PROC QLEARN statement.

5.3 Options for Input

- **DATA**

This option indicates the data set to be used by this PROC.

- **CONTRAST1**

This option indicates the contrast matrix that is used for the first-stage regression. This contrast matrix can be used to obtain estimates for different linear combinations of the

estimated parameters in the stage 1 regression model. Typically, users will specify linear combinations that correspond to differences in the mean of the stage 1 dependent variables for different subgroup of participants based on baseline covariates. The number of columns in this matrix should be equal to number of parameters in the stage 1 regression. The number of rows should correspond to the number of contrasts the user wishes to estimate (see Section 3.2). The default is the identity matrix.

User Tip: The ordering of the columns of the contrast matrix must match the ordering of the parameters in the stage 1 regression model. See Section 3.2 and the user tip at the end of Section 5.5 for more detail.

User Tip: The contrast option can only be used for the stage 1 regression model. To estimate the contrasts for the stage 2 model, PROC GLM can be used.

- **DERIVECI**

This option selects whether or not the confidence intervals for the stage 1 model will be generated. The PROC will generate the confidence intervals in the presence of this statement.

User Tip: To obtain the confidence intervals for the stage 2 regression parameters, PROC GLM can be used.

5.4 Required Statements for PROC QLEARN

The following statements are required by PROC QLEARN in order to specify the regression models for each stage:

- **TAILOR1** *variables ;*

This statement lists the variables for the stage 1 regression model that interact with the first-stage treatment (candidate tailoring variables). This PROC automatically includes "1" as the last element in this statement. In section 3, these variables are the variables in H_{12}. If no variables are listed then the PROC only includes a "1."

- **TAILOR2** *variables ;*

This statement lists the variables for the stage 2 regression model that interact with the second-stage treatment (candidate tailoring variables). This PROC automatically includes "1" as the last element in this statement. In section 3, these variables are the variables in H_{22}. If no variables are listed then the PROC only includes a "1."

- **RESPONSE** *variable ;*

This statement is used to specify the response at the end of the second-stage (final outcome). Note that the final outcome is assumed to be observed for everyone. This variable must be continuous. In section 3, this variable is denoted by Y.

- **STG1TRT** *variable ;*

This statement is used to specify the treatment variable for the stage 1 regression model. This variable must be binary and coded as -1 and +1. In section 3, this variable is denoted A_1.

- **STG2TRT** *variable ;*

This statement is used to specify the treatment variable for the stage 2 regression model. This variable must be binary and coded as -1 and +1. In section 3, this variable is denoted A_2.

5.5 Optional Statements for PROC QLEARN

- **MAIN1** *variables ;*

This statement lists the variables for the stage 1 regression model that do not interact with the first-stage treatment. This PROC automatically includes "1" as the first element in this statement. In section 3, these variables are the variables in H_{11} that are not in H_{12}.

- **MAIN2** *variables ;*

This statement lists the variables for the stage 2 regression model that do not interact with the second-stage treatment. This PROC automatically includes "1" as the first element in this statement. In section 3, these variables are the variables in H_{21} that are not in H_{22}.

- **STG2SAMPLE** *variables ;*

This binary variable indicates whether a participant was rerandomized or not at the stage 2. This variable specifies which observations will be used in the stage 2 regression model. Default = 1, which means all the participants will be included in the stage 2 regression if the statement is omitted. In section 3, this variable is denoted by S.

- **NBOOTSTRAP1** *value ;*

This statement specifies the number of bootstrap samples used to calculate the confidence intervals for the parameters in stage 1. Valid values are integers greater than or equal to 1. Default = 1000.

- **NGRID** *value ;*

This statement specifies the number of grid points over which the algorithm searches to form the upper and lower bounds of the confidence interval for each parameter. Default = 10.

- **GRID-SCALE** *value ;*

This statement specifies the bounds of the grid search for each parameter. Default = 2.58.

- **ALPHA** *value;*
  This statement specifies significance value of the confidence intervals for first-stage parameters and contrasts. Default = 0.05.

Note: In this PROC, variables which are listed in TAILOR1 (TAILOR2) will be added automatically to MAIN1 (MAIN2). Therefore, if a variable appears in both TAILOR and MAIN statements, the PROC will generate an error massage due to multicollinearity.

Note: This PROC adds an intercept to the first- and second-stage models by default so the column "1" does not need to be added to the list of variables specified in MAIN or TAILOR statements.

User Tip: The PROC QLEARN automatically puts the variables in the stage 1 regression parameters in the following order:

- Stage 1: Intercept+ MAIN1+TAILOR1+(TAILOR1+1) STG1TRT

The columns of the contrast matrix discussed in Sections 3.2 and 5.3 must match this ordering.

Table 1: Summary of PROC QLEARN Syntax

Statement	Description
PROC QLEARN	Invokes the procedure
MAIN1	Variables at stage 1 that do not interact with the treatment
TAILOR1	Variables at stage 1 that interact with the treatment (candidate tailoring variables)
MAIN2	Variables at stage 2 that do not interact with the treatment
TAILOR2	Variables at stage 2 that interact with the treatment (candidate tailoring variables)
RESPONSE	The continuous outcome variable
STG2SAMPLE	Rerandomization status as stage 2.
	1 = randomized 0 = not randomized. (default = 1)
STG1TRT	Binary treatment variable at stage 1. Coded -1 and +1.
STG2TRT	Binary treatment variable at stage 2. Coded -1 and +1.
NBOOTSTRAP1	Number of bootstraps to be used in stage 1
	(default = 1000)
NGRID	Number of grid points (default = 10)
GRID-SCALE	Scale for grid range (default = 2.58)
ALPHA	Significance value of the confidence intervals (default = 0.05)

References

Bickel, P. & Freedman, D. (1981). Some asymptotic theory for the bootstrap. *The Annals of Statistics, 9*(6), 1196–1217.

Chakraborty, B., Murphy, S. A., & Strecher, V. (2010). Inference for non-regular parameters in optimal dynamic treatment regimes. *Statistical Methods in Medical Research 19*(3), 317–343.

Kasari, C. (2009). Developmental and augmented intervention for facilitating expressive language, Available at http://clinicaltrials.gov/ct2/show/NCT01013545?term =kasari&rank=5.

Kasari, C., Freeman, S., & Paparella, T. (2006). Joint attention and symbolic play in young children with autism: a randomized controlled intervention study. *Journal of Child Psychology and Psychiatry, 47(*6): 611–620.

Kasari, C., Paparella, T., Freeman, S., & Jahromi, L. B. (2008). Language outcome in autism: Randomized comparison of joint attention and play interventions. *Journal of Consulting and Clinical Psychology, 76*(1):125.

Laber, E., & Murphy, S. (2012). Adaptive confidence intervals for the test error in classification. *Journal of the American Statistical Association.* Advance online publication. doi: 10.1198/jasa.2010.tm10053

Laber, E., Qian, M., Lizotte, D., & Murphy, S. (2010). Statistical inference in dynamic treatment regimes. *Arxiv preprint arXiv:*1006.5831.

Lavori, P., Dawson, R., & Rush, A. (2000). Flexible treatment strategies in chronic disease: Clinical and research implications. *Biological Psychiatry 48*(6), 605–614.

Lei, H., Nahum-Shani, I., Lynch, K., Oslin, D., & Murphy, S. (2011). A" SMART" design for building individualized treatment sequences. *Annual review of clinical psychology*. Advance online publication. doi: 10.1146/annurev-clinpsy-032511-143152

Lunceford, J., Davidian, M., & Tsiatis, A. (2002). Estimation of survival distributions of treatment policies in two-stage randomization designs in clinical trials. *Biometrics 58*(1), 48–57.

Murphy, S. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society, Series B, 65*(2), 331–355.

Murphy, S. (2005). An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine, 24*(10), 1455–1481.

Murphy, S., Van Der Laan, M., & Robins, J. (2001). Marginal mean models for dynamic regimes. *Journal of the American Statistical Association, 96*(456), 1410–1423.

Nahum-Shani, I., Qian, M., Almirall, D., Pelham, W., Gnagy, B., Fabiano, G., Waxmonsky, J. Yu, J., & Murphy, S. (2010). *Experimental design and primary data analysis methods for comparing adaptive interventions* (Technical Report No. 10-108). University Park: The Methodology Center, Penn State. Retrieved from methodology.psu.edu

Nahum-Shani, I., Qian, M., Pelham, W., Gnagy, B., Fabiano, G., Waxmonsky, J., Yu, J., & Murphy, S. (2010). *Q-learning: A data analysis method for constructing adaptive interventions (*Technical Report No. 10-107). University Park: The Methodology Center, Penn State. Retrieved from methodology.psu.edu

Pelham, W., & Fabiano, G. (2008). Evidence-based psychosocial treatments for attention-deficit/hyperactivity disorder. *Journal of Clinical Child & Adolescent Psychology, 37*(1), 184–214.

Pelham, W. E., Hoza, B., Pillow, D. R., Gnagy, E.M., Kipp, H. L., Greiner, A. R. , Waschbusch, D. A., Trane, S. T., Greenhouse, J., Wolfson, L. (2002). Effects of methyphenidate and expectancy on children with adhd: Behavior, academic performance, and attributions in a summer treatment program and regular classroom settings. *Journal of Consulting and Clinical Psychology, 70*(2): 320.

Robins, J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period–application to control of the healthy worker survivor effect. *Mathematical Modelling, 7*(9-12), 1393–1512.

Robins, J. (2004). Optimal structural nested models for optimal sequential decisions. In *Proceedings of the Second Seattle Symposium on Biostatistics* (pp. 189–326). New York, NY: Springer.

Thall, P., Sung, H., & Estey, E. (2002). Selecting therapeutic strategies based on efficacy and death in multicourse clinical trials. *Journal of the American Statistical Association, 97*(457), 29–39.

Thall, P. & Wathen, J. (2005). Covariate-adjusted adaptive randomization in a sarcoma trial with multi-stage treatments. *Statistics in medicine 24*(13), 1947–1964.

Tsiatis, A. (2006). *Semiparametric theory and missing data*. New York, NY: Springer.

Wahed, A., & Tsiatis, A. (2004). Optimal estimator for the survival distribution and related quantities for treatment policies in two-stage randomization designs in clinical trials. *Biometrics, 60*(1), 124–133.

Wahed, A., & Tsiatis, A. (2006). Semiparametric efficient estimation of survival distributions in two-stage randomisation designs in clinical trials with censored data. *Biometrika, 93*(1), 163–177.

Watkins, C. J. C. H. (1989). *Learning from delayed rewards* (Doctoral thesis, University of London, King's College. Available from ProQuest Dissertations and Theses Database. (ProQuest Document ID No. 303776498.

Watkins, C. J. C. H. & Dayan, P. (1992). Q-learning. *Machine Learning, 8*(3), 279–292.

Appendices

In the Appendices, we apply PROC QLEARN to two simulated data sets. These simulated data sets mimic data arising from two different SMART studies: (i) a study to develop an adaptive treatment strategy for children with ADHD (for more details see Pelham et al., 2002; Pelham & Fabiano, 2008; and Nahum-Shani et. al. 2010b), and (ii) a study to develop an adaptive treatment strategy for minimally verbal children with autism spectrum disorder (for details see Kasari, 2009; Kasari et al., 2006; Kasari et al., 2008; and Lei et al., 2011;).

Appendix A - Adaptive Interventions for Children with ADHD

Attention-deficit hyperactivity disorder (ADHD) is a chronic disorder among children which is characterized by trouble controlling impulsive behavior. There are a number of pharmacological and behavioral modification treatments available for treating children with ADHD. It is not clear, however, whether the stage 1 treatment should be medication or behavioral modification. In addition, it is not clear how to treat children who do not respond to the initial treatment offered (either medication or behavioral modification). For example, for children who do not respond to the stage 1 treatment (e.g., children who do not respond to their medicine), is it better to intensify the stage 1 treatment or to augment with a new treatment?

An example ADHD SMART involving 150 simulated children is shown in Figure 1. The SMART has the following features:

- At stage 1, all children are randomized to either behavioral modification (coded as a_1 =+1) or to medication (coded as a_1=-1).
- Beginning at 8 weeks (two months), the child's response to treatment is assessed monthly until the end of the school year (for a total of 8 monthly assessments). Based on their monthly scores, the children are labeled as responders or non-responders to the first-stage treatment. Children who are labeled non-responders are moved to stage 2 treatment; therefore, note that the amount of time spent in stage 1 treatment may vary from child to child.

- At stage 2, non-responders to low-intensity behavioral modification are rerandomized to either intensified behavioral modification (coded as a_2 =+1) or to behavioral modification augmented with medication (coded as a_2 =-1). Non-responders to medication are rerandomized to either an increased dose of medication (coded as a_2 =+1) or to medication augmented with behavioral modification (coded as a_2=-1). 99 children are non-responders in the simulated data set. Responders continue their stage 1 treatment.
- After eight months, the school performance, y, of each child is assessed.

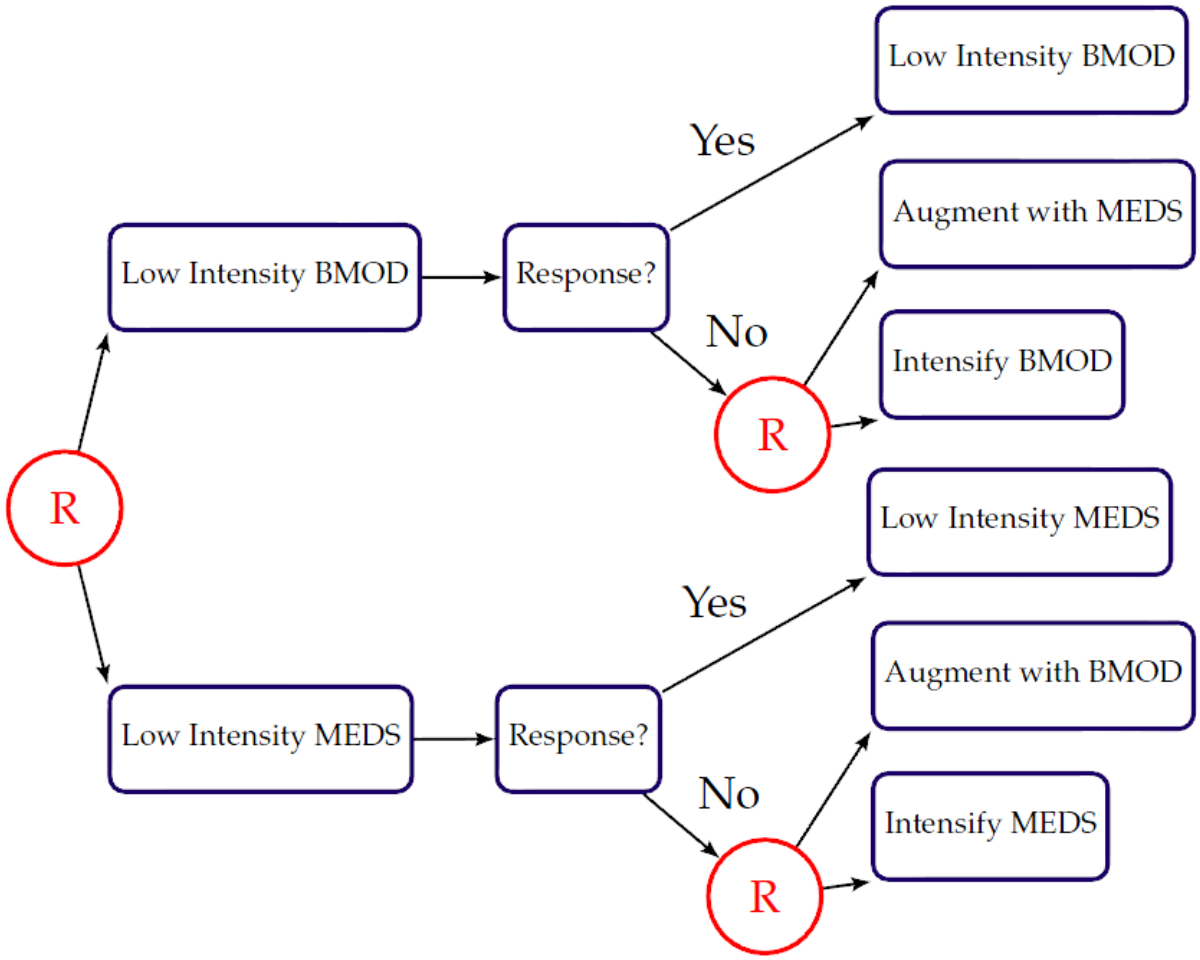

Figure 1: SMART design of the ADHD study. BMOD, behavior modification; MEDS, medication. The "R" with a circle around it denotes randomization.

The simulated data set is called `adhddat.` In addition to the treatment variables a_1 and a_2, it includes a number of measures that could serve as baseline tailoring variables for making a stage 1 decision between medication vs behavioral modification. The candidate baseline tailoring variables include:

- o_{11}: ODD Diagnosis: whether the child was (coded as 1) or was not (coded as 0) diagnosed with ODD (oppositional defiant disorder) before the first-stage treatment.
- o_{12}: ADHD symptoms at the end of the previous school year (ranging from 0 to 3, larger values reflect fewer symptoms).
- o_{13}: Medication prior to first-stage treatment, reflecting whether the child did (coded as 1) or did not (coded as 0) receive medication during the previous school year.
- o_{14}: Race variable: white (coded 1) versus non-white (coded 0).

The data set also includes candidate time-varying tailoring variables from stage 1. These variables can be used, along with the candidate baseline tailoring variables, to make a decision between intensifying vs augmenting treatment among those who are labeled non-responders to stage 1 treatment. The candidate time-varying tailoring variables include:

- o_{21}: Month during the school year at which the child showed inadequate response to the first-stage treatment, and hence entered the second-stage of the treatment.
- o_{22}: Adherence to first-stage treatment: whether the child did (coded as 1) or did not (coded as 0) show high adherence to the first-stage treatment.

The simulated data set includes two additional variables, r and s, which are the responder and rerandomization indicators, respectively. r is coded 1 for the responders and 0 for non-responders. Since all non-responders (and no responders) to the first-stage treatment are rerandomized at the second-stage, only the sample of 99 non-responders can be used in the stage 2 regression model. As a result, s is coded 1 if the child is non-responder and is coded 0 otherwise. As such, in this example, $s = 1 - r$. Below, we show the first 6 observations of the simulated data set.

ID	o11	o12	o13	o14	a1	r	o21	o22	a2	y	s
5	0	-0.70404	1	0	-1	1	.	0	.	4	0
9	1	-1.01254	1	1	-1	1	.	0	.	4	0
23	0	0.14336	1	1	-1	1	.	1	.	5	0
24	1	1.07187	0	1	-1	1	.	0	.	1	0
27	0	-0.84515	0	0	-1	1	.	1	.	4	0
30	0	-0.21032	0	1	1	0	7	0	1	1	1

By design, the variables o_{21} and a_2 are relevant only for the non-responders ($r = 0$) of the stage 1 treatment. Therefore, they are denoted by "." for those who are responders ($r = 1$) in the dataset.

Below, the covariates are centered and are denoted with a subscript c. This is not necessary, but it facilitates interpretation of the regression parameters.

```
* center candidate baseline and time-varying tailoring variables; data dat1;
    set adhddat;
    o11c = o11 -  0.3533333;
    o12c = o12 - -0.1205948;
    o13c = o13 -  0.3133333;
    o14c = o14 -  0.8066667;
    o21c = o21 -  4.5858586;
    o22c = o22 -  0.4533333;
run;
```

Suppose we decide to fit the following regression models:

- Stage 1: $\beta_{10} + \beta_{11}o_{11c} + \beta_{12}o_{12c} + \beta_{13}o_{13} + \beta_{14}a_1o_{13} + \beta_{15}a_1$
- Stage 2: $\beta_{20} + \beta_{21}o_{11c} + \beta_{22}o_{12c} + \beta_{23}o_{13c} + \beta_{24}o_{14c} + \beta_{25}o_{21c} + \beta_{26}a_1 + \beta_{27}o_{22} + \beta_{28}a_1a_2 + \beta_{29}a_2o_{22} + \beta_{210}a_2$

For illustrative purposes, the model at stage 1 has only one interaction term involving first stage treatment (a_1); namely, the β_{14} interaction term between the baseline variable measuring medication in the prior year o_{13} and first stage treatment a_1. Thus, o_{13} is the only candidate baseline tailoring variable considered in this example regression model. In order to examine whether baseline prior year medication is a tailoring variable, we can use the following contrast matrix in PROC QLEARN.

```
data contrast1;
    input M1 M2 M3 M4 M5 M6;
datalines;
 1 0 0 1 1 1 /* mean outcome under BMOD for children with o13=1*/
 1 0 0 1-1-1 /* mean outcome under MED for children with o13=1*/
 0 0 0 0 2 2 /* mean difference BMOD-MED for children with o13=1*/
 1 0 0 0 0 1 /* mean outcome under BMOD for children with o13=0*/
 1 0 0 0 0-1 /* mean outcome under MED for children with o13=0*/
 0 0 0 0 0 2 /* mean difference BMOD-MED for children with o13=0*/
;
run;
```

Note that the contrast matrix has 6 columns corresponding to the fact that there are 6 parameters in the stage 1 regression. The number of rows is determined by the user; it corresponds to the number of linear contrasts the user wishes to estimate. In the example above, we explain the meaning of each linear contrast in comments to the right of each row.

The following code presents the PROC QLEARN syntax to fit the above stage 1 and stage 2 regression models.

```
proc qlearn data=dat1 contrast1=contrasts1 deriveci;
   main1 o11c o12c;
   tailor1 o13;
   main2 o11c o12c o13c o14c o21c;
   tailor2 a1 o22;
   response y;
   stg2sample s;
   stg1trt a1;
   stg2trt a2;
run;
```

Results:

Table 2 summarizes the results obtained by PROC QLEARN on this simulated data. The first contrast estimates the mean outcome (3.19) among children who are assigned to behavioral modification ($a_1 = 1$) at stage 1 and who received medication in the prior year $(o_{13} = 1)$; o_{11c} and o_{12c} were set to 0. The second contrast estimates the mean outcome (3.66) among children who are assigned to medication ($a_1 = -1$) at stage 1 and who received medication in the prior year $(o_{13} = 1)$; o_{11c} and o_{12c} were set to 0. The third contrast estimates the effect of behavioral modification ($a_1 = 1$) vs medication ($a_1 = -1$) on the mean outcome (-0.46) when $o_{13} = 1$.

Contrast 4 estimates the mean outcome (3.74) among children who are assigned to behavioral modification ($a_1 = 1$) at stage 1 and who did not receive medication in the prior year $(o_{13} = 0)$; o_{11c} and o_{12c} were set to 0. Contrast 5 estimates the mean outcome (3.16) among children who are assigned to medication ($a_1 = -1$) at stage 1 and who did not receive medication in the prior year $(o_{13} = 0)$; o_{11c} and o_{12c} were set to 0. The last contrast estimates the effect of behavioral modification ($a_1 = 1$) vs medication ($a_1 = -1$) on the mean outcome (0.59) when $o_{13} = 0$.

The estimated mean outcome using contrast 1 & 2 suggests that children who received medication in the prior year (i.e. $o_{13} = 1$) benefit from medication, and Contrast 3 shows that the medication works significantly better on this group of children since the 95% confidence interval does not include 0. The estimated mean outcome using contrast 4 & 5 suggests that children who did not receive medication in the past (i.e. $o_{13} = 0$) benefit from behavioral modification, and Contrast 6 shows that on this group of children, the behavioral modification works significantly better since the 95% confidence interval does not include 0.

Table 2. ADHD study PROC QLEARN output.

Number of observations in dataset: 150
Number of observations in stage 1: 150
Number of observations in stage 2: 99

First Stage Regression Result

Variable	Parameter Estimates	Confidence Interval Upper	Lower
intercept	3.4497	3.6621	3.0956
o11c	-0.4556	-0.0660	-0.8855
o12c	-0.3458	-0.1585	-0.5136
o13	-0.0236	0.3938	-0.4827
o13 :a1	-0.5254	-0.2418	-0.7742
a1	0.2934	0.4570	0.1008

Contrasts	Parameter Estimates	Confidence Interval Upper	Lower
Contrast 1	3.1941	3.5891	2.6733
Contrast 2	3.6580	4.0424	3.3163
Contrast 3	-0.4639	-0.0362	-0.9799
Contrast 4	3.7431	4.0303	3.3799
Contrast 5	3.1563	3.4306	2.7876
Contrast 6	0.5868	0.9140	0.2016

Second Stage Regression Result

Variable	Parameter Estimates
intercept	3.0039
o11c	-0.2462
o12c	-0.2961
o13c	0.0391
o14c	0.4868
o21c	-0.0097
a1	0.0758
o22	-0.0980
a1 :a2	-0.1934
o22 :a2	1.1826
a2	-0.8640

Appendix B: Adaptive Interventions for Minimally Verbal Children

Minimally verbal children (i.e. children who communicate with fewer than 20 functional words) with autism spectrum disorder often do not respond to traditional intensive treatments. Interventions for these children require an intervention strategy that incorporates a plan for what treatment to provide if the child does not initially respond.

The example Autism SMART involving 200 simulated children is shown in Figure 2. The design has the following features:

- At stage 1, all children are randomized to either the joint attention/joint engagement (JAE) treatment combined with the enhanced milieu teaching (EMT) treatment (coded as $a_1 = 1$) or the joint attention/joint engagement (JAE) treatment combined with the augmentative and alternative communication (AAC) treatment (coded as $a_1 = -1$).
- After 3 months, all children are classified as responders or non-responders to the stage 1 treatment based on their improvement in spoken communication.
- At stage 2, responders continue with their initial treatment for an extra 3 months. Children who did not respond to JAE+AAC continue with intensified JAE+AAC. Non-responders to JAE+EMT, however, are randomized to either JAE+AAC (coded as $a_2 = -1$) or intensified JAE+EMT (coded as $a_2 = +1$). 56 children in the simulated data set do not respond to JAE+EMT.
- After 6 months, the number of different spontaneous words is assessed (y).

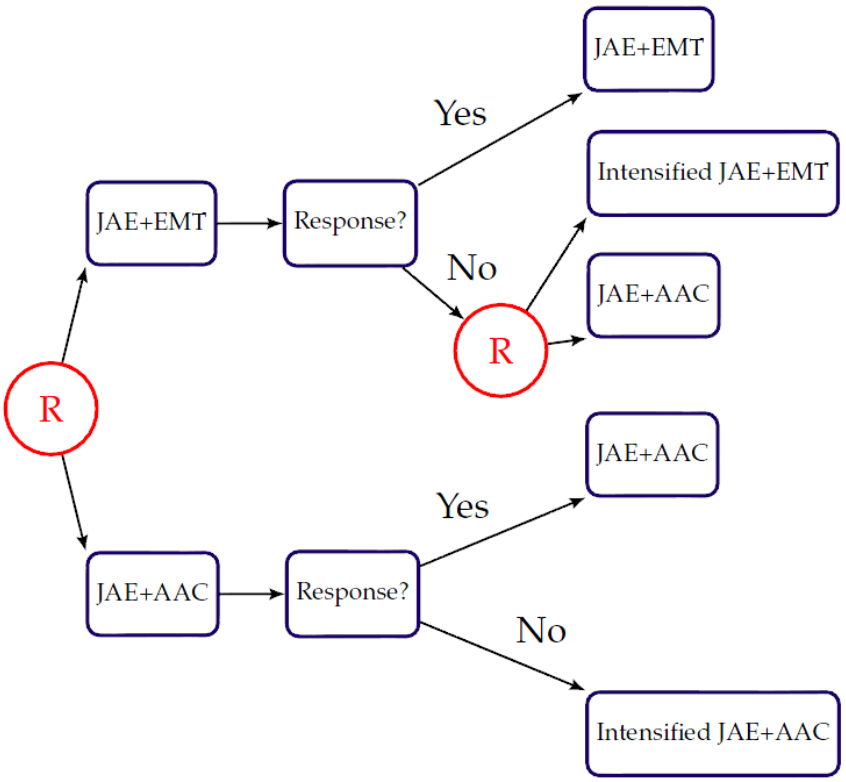

Figure 2: SMART design of the Autism study. AAC, augmentative and alternative communication; EMT, enhanced milieu teaching; JAE, joint attention/joint engagement. The "R" with a circle around it denotes randomization.

The simulated data set is called autismdat. It contains a number of measures that can serve as baseline tailoring variables for making a stage 1 decision between JAE+EMT vs JAE+AAC. The candidate baseline tailoring variables include:

- o_{11}: Number of different spontaneous words ($o_{11} \leq$ 10: low and $o_{11} \geq$ 70: high). High values are preferred. This is the same measure as the outcome y, but is measured at baseline.
- o_{12}: Number of unintelligible utterances by the child. Lower values are preferred.

The data set also contains candidate time-varying tailoring variables from stage 1. These variables can be used, along with the candidate baseline tailoring variables, to make a decision between intensifying vs augmenting treatment among those who are labeled as non-responders to initial JAE+EMT. The candidate time-varying tailoring variables include:

- o_{21}: Number of communicative functions by the child (e.g., to request, to initiate a social interaction, to direct another's attention to an object) during stage 1 treatment. Higher values are preferred.
- o_{22}: Number of spontaneously communicative utterances by the child (i.e., the number of times the child spoke unprompted). This does not include words used to imitate or if the child is prompted to speak. Higher values are preferred.

The simulated data set includes one additional variable, s, which is the rerandomization indicator. Since only non-responders to the first-stage treatment with JAE+EMT are rerandomized at the second-stage, only these children's data can be used in the stage 2 regression model. As a result s is coded 1 if the child is non-responder to JAE+EMT and is coded 0 otherwise.

Suppose we decide to fit the following regression models:

- Stage 1: $\beta_{10} + \beta_{11}o_{11} + \beta_{12}a_1o_{11} + \beta_{13}a_1$
- Stage 2: $\beta_{20} + \beta_{21}o_{11} + \beta_{22}o_{21} + \beta_{23}a_2o_{21} + \beta_{24}a_2$

The model at stage 1 has only one interaction term involving first stage treatment (a_1) (between baseline number of different spontaneous words (o_{11}) and a_1); thus, in this illustrative analysis o_{11} is the only candidate baseline tailoring variable. In order to examine whether baseline number of different spontaneous words is a tailoring variable, we use the following contrast matrix in PROC QLEARN.

```
data contrasts1;
  input C1 C2 C3 C4;
  datalines;
  1  10  10   1  /*mean outcome under EMT for children w/o11 = 10 = low;*/
  1  10 -10  -1  /*mean outcome under AAC for children w/o11 = 10 = low;*/
  0   0  20   2  /*mean difference EMT - AAC for children w/o11 = 10 = low;*/
  1  70  70   1  /*mean outcome under EMT for children w/o11 = 70 = hi;*/
  1  70 -70  -1  /*mean outcome under AAC for children w/o11 = 70 = hi;*/
  0   0 140   2  /*mean difference EMT - AAC for children w/o11 = 70 = hi;*/
  ;
  run;
```

Note that the contrast matrix has 4 columns corresponding to the fact that there are 4 parameters in the stage 1 regression.

The following code presents the PROC QLEARN syntax to fit the above stage 1 and 2 regression models.

```
proc qlearn data=autismdat contrast1=contrasts1 deriveci;
  tailor1 o11;
  main2 o11;
  tailor2 o21;
  response y; s
  tg2sample s;
  stg1trt a1;
  stg2trt a2;
run;
```

Results:

Table 3 summarizes the results obtained by the PROC QLEARN on this simulated data. The first contrast estimates the mean outcome (51.37) under EMT ($a_1 = 1$) for children when $o_{11} = 10$ (e.g., baseline number of spontaneous words, o_{11}, is low). The second contrast estimates the mean outcome (57.51) under AAC ($a_1 = -1$) for children when $o_{11} = 10$. The third contrast estimates the effect of JAE+EMT ($a_1 = 1$) vs JAE+AAC ($a_1 = -1$) on the mean outcome (-6.14) when the baseline number of spontaneous words, o_{11}, is low.

Contrast 4 estimates the mean outcome (82.46) under EMT ($a_1 = 1$) for children when $o_{11} = 70$ (i.e. baseline number of spontaneous words, o_{11}, is high). Contrast 5 estimates the mean outcome (60.88) under AAC ($a_1 = -1$) for children when $o_{11} = 70$. The last contrast estimates the effect of JAE+EMT ($a_1 = 1$) vs JAE+AAC ($a_1 = -1$) on the mean outcome (21.57) when the baseline number of spontaneous words, o_{11}, is high.

The estimated mean outcome using contrast 1 & 2 suggests that patients with low baseline number of spontaneous words, $o_{11} = 10$, benefit from JAE+AAC; however, contrast 3 shows that this difference is not significant since the 95% confidence interval includes zero. The

estimated mean outcome using contrast 4 & 5 suggests that patients with high baseline number of spontaneous words, $o_{11} = 70$, benefit from JAE+EMT and contrast 6 confirms that the difference is significant since the 95% confidence interval does not include zero.

Table 3. Autism study PROC QLEARN output.

Number of observations in dataset:	200
Number of observations in stage 1:	200
Number of observations in stage 2:	56

First Stage Regression Result

Variable	Parameter Estimates	Confidence Interval Upper	Lower
intercept	51.5675	55.9073	45.2093
o11	0.2872	0.4367	0.1446
o11 :a1	0.2309	0.3570	0.1281
a1	-5.3776	-1.1479	-9.6248

Contrasts	Parameter Estimates	Confidence Interval Upper	Lower
Contrast 1	51.3710	55.5031	47.8922
Contrast 2	57.5077	63.1635	51.2595
Contrast 3	-6.1367	0.4455	-15.1740
Contrast 4	82.4582	86.5363	77.4660
Contrast 5	60.8831	67.1254	50.1310
Contrast 6	21.5750	31.4390	13.1934

Second Stage Regression Result

X Variable	Parameter Estimates
intercept	28.3800
o11	0.5229
o21	6.2355
o21 :a2	6.2175
a2	-26.8556