Analytic and Sample Size Considerations for Sequential, Multiple-Assignment Randomized Trials with Longitudinal Outcomes

Thesis Proposal

by

Nicholas J. Seewald

Statistics
University of Michigan
2018

Doctoral Committee:

Research Associate Professor Daniel Almirall, Co-Chair
Professor Kerby Shedden, Co-Chair
Professor Naisyin Wang
Research Associate Professor Kelley M. Kidwell

2018

TABLE OF CONTENTS

List of Figures . . . . . **ii**

List of Tables . . . . . **iii**

Chapter

1 An Introduction to Dynamic Treatment Regimens and SMART Designs **1**

1.1 Dynamic Treatment Regimens 2
1.2 Sequential Multiple-Assignment Randomized Trials 3

2 SMARTs with Continuous Repeated-Measures Outcomes **8**

2.1 Marginal Mean Model 8
2.2 Estimation 11
2.2.1 Observed Data 11
2.2.2 Estimating Equations 11
2.2.3 Estimation of the Working Covariance Matrix 13
2.2.4 Iterated Estimation Procedure 14
2.3 Sample Size Formulae for End-of-Study Comparisons of Embedded DTRs in Two-Stage SMARTs 14
2.4 Simulation Results 18
2.5 Discussion 22

3 Future Work **24**

3.1 Next Steps for Sample Size Calculation in Repeated-Measures SMARTs . 24
3.1.1 Work Remaining Prior to Manuscript Submission 24
3.1.2 Software Tools for SMARTs with Continuous Longitudinal Outcomes 25
3.2 Sample Size for Repeated-Measures SMARTs with Budget Constraints . 26
3.3 SMARTs with Intensive Longitudinal Outcomes 27

Appendices **29**

Bibliography **37**

LIST OF FIGURES

1.1 Three commonly-used two-stage SMART designs 5
1.2 The ENGAGE SMART 7

2.1 Plot of example marginal structural mean model for ENGAGE 10
2.2 Empirical power versus the difference between the true within-person correlation ρ and hypothesized correlation ρ_{guess} used to compute sample size 20

LIST OF TABLES

2.1 Design-specific indicators for consistency with a given DTR 12
2.2 Correlation estimators for selected working correlation structures 14
2.3 SMART-specific design effects 18
2.4 Sample sizes and empirical power results for an end-of-study comparison of the DTR recommending only treatments indexed by 1 and that which recommends only treatments indicated by -1 21

A.1 Design-specific consistency assumptions 30

CHAPTER 1

An Introduction to Dynamic Treatment Regimens and SMART Designs

In practice, interventions often involve sequences of treatments that are adapted to an individual's changing needs. A single, fixed treatment may or may not be adequately effective for all individuals at all times; indeed, heterogeneity of treatment effects across people often exists (Longford 1999; Gail and Simon 1985). Chronic conditions which wax and wane in severity may require an intervention strategy which adjusts treatment according to changing severity over time.

Clinical practice typically involves the provision of treatment, some follow-up period, then modification of treatment to better suit the individual's needs, if necessary. However, open questions often remain as to the protocolization of this sequence. For example, "[i]gnorance of whether or how to change psychotherapies is a major and persisting gap in psychiatric knowledge" (Markowitz and Milrod 2015).

Dynamic treatment regimens (DTRs) operationalize clinical decision-making by recommending particular treatments to certain subsets of patients at specific times (Chakraborty and Moodie 2013). DTRs are sequences of pre-specified decision rules leading to courses of treatment which adapt to a patient's changing needs (Kosorok and Moodie 2015). Consider the following example DTR which was designed to increase engagement with an intensive outpatient rehabilitation program (IOP) for patients with alcohol and/or cocaine dependence: "Within a week of the participant becoming non-engaged in the IOP, provide a phone-based session focused on helping the patient re-engage in the IOP. At week 8, look back at the participant's engagement pattern over the past eight weeks. If the participant continued to not engage, provide a second phone-based session, this time focused on facilitating personal choice (i.e., highlighting various treatment options the patient can choose from in addition to IOP). Otherwise, provide no further contact" (McKay et al. 2015). Notice that the DTR recommends intervention strategies for both engaged and non-engaged participants at week 8. Alternative names for DTRs include adaptive treatment strategies

(Wallace and Moodie 2014; Ogbagaber, Karp, and Wahed 2016) and adaptive interventions (Almirall et al. 2014; Nahum-Shani et al. 2012a), among others.

Scientists often have questions about how best to sequence and individualize interventions in the context of a DTR. Sequential, multiple-assignment, randomized trials (SMARTs) are one type of randomized trial design that can be used to answer questions at multiple stages of the development of high-quality DTRs (Lavori and Dawson 2000, 2004; Murphy 2005). The characteristic feature of a SMART is that some or all participants are randomized more than once, often based on previously-observed covariates. Each randomization corresponds to a critical question regarding the development of a high-quality DTR, typically related to the type, timing, or intensity of treatment. SMARTs have been employed in a variety of fields, including oncology (Auyeung et al. 2009; Kidwell 2014; Thall 2015), surgery (Diegidio et al. 2017; Hibbard et al. 2018), substance abuse (Murphy et al. 2007), and autism (Kasari et al. 2014).

In this chapter, we introduce and motivate the study of DTRs and SMARTs. We begin by formally defining DTRs, then discuss how they can be studied using a SMART. We present a variety of SMART designs, and discuss motivations for each.

1.1 Dynamic Treatment Regimens

A DTR is a sequence of functions ("decision rules"), each of which takes as inputs a person's history up to the time of the current decision (including baseline covariates, adherence, responses to previous treatments, etc.) and outputs a recommendation for the next treatment (Murphy 2005). Formally, suppose we wish to construct a DTR which recommends M treatments, $a_1, \ldots, a_M$, to each individual. Let S_j denote information collected in the period after providing treatment a_{j-1} until immediately prior to the provision of treatment a_j, and let $\bar{a}_j = \{a_1, a_2, \ldots, a_j\}$ denote the sequence of treatments provided up to and including treatment a_j. Define $\bar{S}_j(\bar{a}_{j-1}) = \{S_1, S_2(a_1), \ldots, S_{j-1}(\bar{a}_{j-2})\}$ to represent the "history" until the time at which a_j is provided. This includes any outcomes and covariates which may be observed, as well as previous treatment assignments. S_1 contains pre-treatment information. Note that $\bar{S}_j(\bar{a}_{j-1})$ is indexed by the history of treatment assignments made up to, but not including, the time at which a_j is assigned, reflecting the fact that different values of the covariates may be observed depending on the assigned sequence of treatments.

A decision rule φ_j is a function of $\bar{S}_j(\bar{a}_{j-1})$ which outputs a recommendation for subsequent treatment a_j. An M-stage dynamic treatment regimen is a sequence of $M-1$ decision rules $\{\varphi_1, \ldots, \varphi_{M-1}\}$ (Murphy 2005).

The information $S_j(\bar{a}_{j-1})$ often contains covariates which inform the recommendation to subsequent treatment a_j. These covariates are called "tailoring variables." Consider the example two-stage DTR above. The recommended first-stage treatment is a phone-based session with a focus on re-engagement with the IOP. At week 8, each participant's history of engagement is assessed, and an appropriate second-stage treatment is recommended. For participants who have shown a pattern of continued non-engagement, the recommended second-stage treatment is a second phone-based session focusing on personal choice. For all other participants, the DTR recommends no further contact. Here, the tailoring variable contained in $S_2(a_1)$ is an indicator as to whether or nor the participant demonstrated a pattern of continued non-engagement prior to week 8.

In the sequel, we will consider only two-stage DTRs, which recommend two treatments separated by a single decision rule. Further, we focus on binary tailoring variables, which we will abbreviate to "response" or "non-response". Since a DTR recommends treatments to both "responders" and "non-responders", we can denote a DTR with a triple of the form (a_1, a_{2R}, a_{2NR}), where a_1 is an indicator for the recommended first-stage treatment, a_{2R} an indicator for the second-stage treatment recommended for responders, and a_{2NR} the second-stage treatment recommended for non-responders.

1.2 Sequential Multiple-Assignment Randomized Trials

A SMART is a type of randomized trial in which some or all participants are randomized more than once. Each randomization in a SMART corresponds to a scientific question regarding whether or how to adapt the type, timing, or intensity of treatment. We consider two-stage SMARTs in which the primary outcome is continuous and repeatedly measured in participants over the course of the study.

Most SMARTs contain "embedded" DTRs; that is, by design, participants in a SMART may be assigned to treatments which are consistent with recommendations made by one or more DTRs. Often, subsequent randomizations in a SMART are restricted to particular groups of participants based on an embedded tailoring variable, which is chosen based on scientific, ethical, or practical considerations. For example, in cancer, it would be unethical to randomize patients who do not respond to a high dose of chemotherapy to an intervention which would increase the dose beyond a known toxicity threshold. Instead, these individuals may not be re-randomized.

We consider SMARTs in which each randomization is between two possible interventions, and where the tailoring variable is binary. Figure 1.1 introduces three common two-stage SMART designs which vary in the subsets of participants who are re-randomized

after the first stage.

In design I, all participants are re-randomized. There are eight DTRs embedded in this design: for example, the DTR which starts by recommending A, then recommends C for responders and F for non-responders. Using the notation above and the indices in figure 1.1, this DTR would be written (1, 1, -1). SMARTs of this form have been run in the fields of drug dependence (Oslin 2005; Fitzsimons et al. 2015), smoking cessation (Fu et al. 2017), and childhood depression (Eckshtain 2013), among others. Often, motivation for re-randomizing all participants arises out of open scientific questions regarding both a maintenance therapy for responders and a "rescue" intervention for non-responders.

SMARTs using design II restrict the second randomization to only non-responders; that is, only participants who have a certain value of the tailoring variable (here, "non-response") are re-randomized. This might be motivated by an open question regarding second-stage treatment only among non-responders (i.e., the follow-up intervention for responders may be well-established). Design II is perhaps the most common SMART design, and it has been utilized in the study of ADHD (Pelham et al. 2016), adolescent marijuana use (Budney 2014), alcohol and cocaine dependence (McKay et al. 2015), and more. There are four embedded DTRs in this design. Because responders are not re-randomized, a_{2R} is set to zero for all embedded DTRs.

In design III, re-randomization is restricted to only non-responders who receive a particular first-stage treatment. This design might be used when one of the first-stage interventions involves a top-of-the-line treatment that, for ethical reasons, cannot be taken away in the second stage. For the individuals randomized to this first-stage treatment, there may only be one option for subsequent intervention, regardless of their value of the tailoring variable. SMARTs of this type have been used to investigate cognition in children with autism spectrum disorder (Kasari et al. 2014; Almirall et al. 2016) and implementation of a re-engagement program for patients with mental illness (Kilbourne et al. 2013). There are three DTRs embedded in this design. Note that, as in design II, responders are not re-randomized, so a_{2R} is set to zero for all embedded DTRs. Furthermore, a_{2NR} is set to zero when $a_1 = -1$, as non-responders to treatment B are not re-randomized.

For more information on various SMART designs and case studies for each type, see Lei et al. (2012).

To illustrate our ideas, we use ENGAGE, illustrated in figure 1.2. ENGAGE is a SMART designed to study the effects of offering cocaine- and/or alcohol-dependent patients who did not engage in an IOP phone-based sessions either geared toward re-engaging them in an IOP or offering a choice of treatment options (McKay et al. 2015). The study recruited 500 cocaine- and/or alcohol-dependent adults who were enrolled in an IOP and

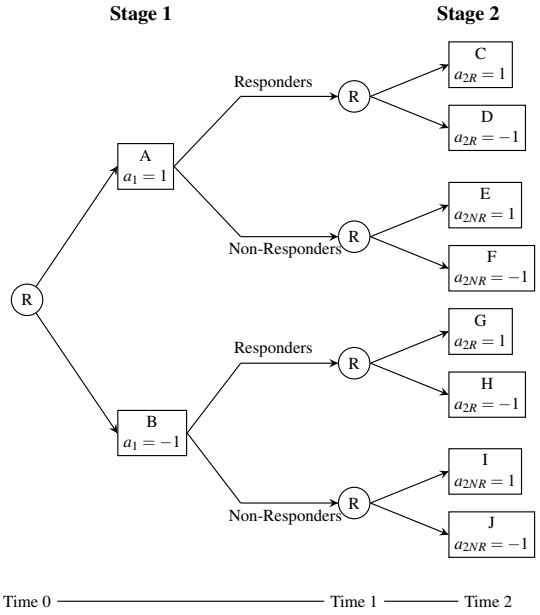

(I) All participants are re-randomized, regardless of response status.

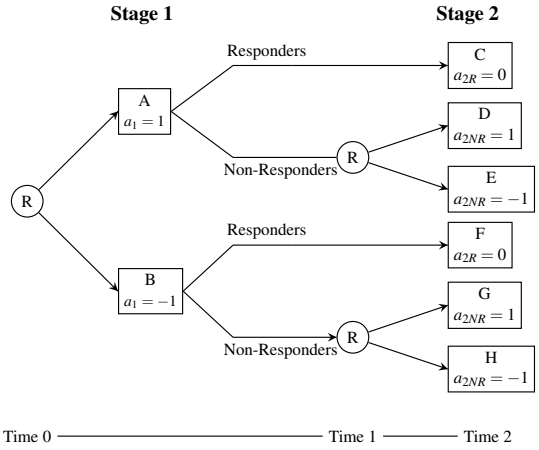

(II) The second randomization is restricted to only non-responders.

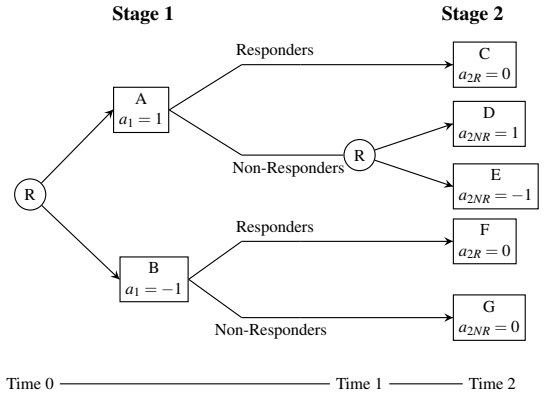

(III) The second randomization is restricted to only non-responders to treatment A.

Figure 1.1: Three commonly-used two-stage SMART designs. Each design varies in choice of which subsets of participants are re-randomized. Circled R indicates randomization, capital letters indicate (potentially non-unique) treatments, and $a_{\_}$ provides a coding system used to index embedded DTRs.

failed to attend two or more sessions in the first two weeks. ENGAGE is modeled on design II. In the context of figure 1.1, treatment A was two phone-based motivational interviews focused on reengaging the participant with the IOP ("MI-IOP"); treatment B was two phone-based motivational interviews geared towards helping the participant choose and engage with an intervention of their choice ("MI-PC"). Participants who exhibited a pattern of continued non-engagement after eight weeks were considered non-responders, and re-randomized to receive either MI-PC (treatments D and G) or no further contact (treatments E and H). Responders were provided no further contact (treatments C and F). Following the coding in figure 1.1, the example DTR on page 1 is labeled (1, 0, 1).

An important continuous outcome in ENGAGE is "treatment readiness". This is a measure of a patient's willingness and ability to commit to active participation in a substance abuse treatment program. The score ranges from 8-40 and is coded so that higher scores indicate greater treatment readiness. Measurements are taken at baseline, and 4, 8, 12, and 24 weeks after program entry.

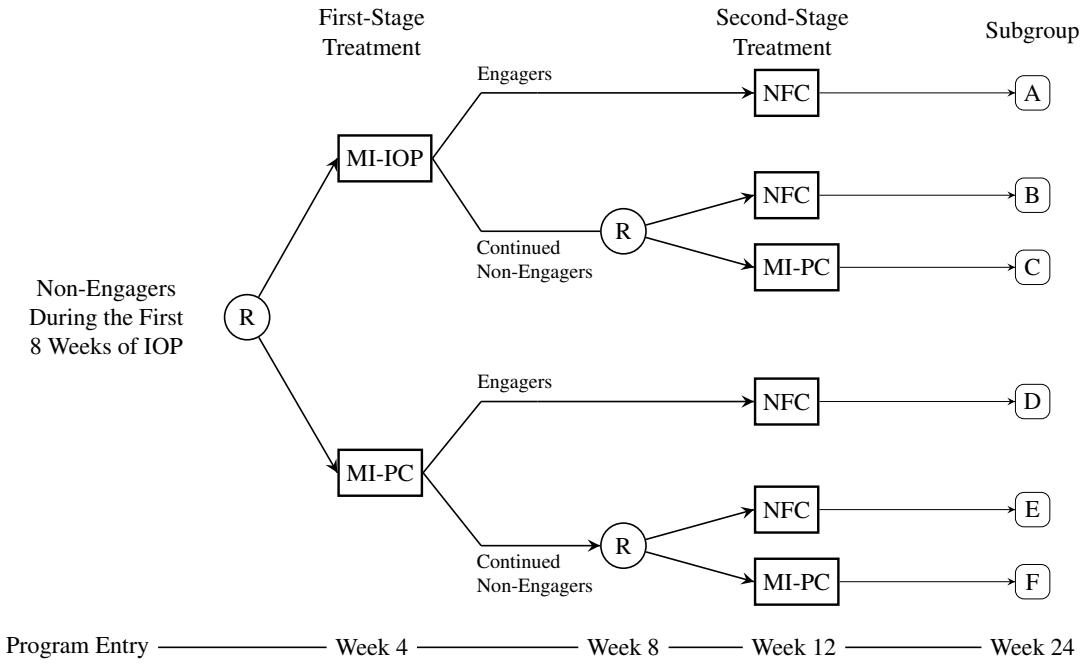

Figure 1.2: The ENGAGE SMART

CHAPTER 2

SMARTs with Continuous Repeated-Measures Outcomes

Often, SMARTs involve repeated measurements of a continuous outcome spaced throughout the trial. This might involve a measurement at baseline, one or more measurements in the first stage of the trial (before assessment of the tailoring variable and subsequent re-randomization), and one or more measurements in the second stage.

In this chapter, we begin by reviewing the work of Lu et al. (2016), in which they developed models and an estimation procedure for SMARTs with longitudinal outcomes. We then extend that work by offering more detailed guidance on the estimation of model parameters used in computing quantities of interest on which to compare two embedded DTRs. Finally, we present sample size formulae for SMARTs in which the primary aim is to compare the mean end-of-study outcomes for two embedded DTRs which recommend different first-stage treatments and which satisfy certain design constraints. Sections 2.2.3 and 2.3 to 2.5 are the primary contribution of a manuscript targeted for *Statistical Methods in Medical Research* (impact factor 2.284 as of 1 December 2018).

2.1 Marginal Mean Model

Consider a SMART design with embedded DTRs labeled by (a_1, a_{2R}, a_{2NR}). Suppose we have a repeated-measures outcome $\boldsymbol{Y} = (Y_{t_1}, \ldots, Y_{t_T})$ observed such that Y_t is measured for all participants at each of T time points $\left\{t_j : j = 1, \ldots, T; t_1 < \ldots < t_T\right\}$. We do not require that these time points be equally-spaced, though they must be common to all participants in the study. Define $t^* \in \left\{t_j\right\}$ to be the measurement occasion immediately before assessment of the tailoring variable and any subsequent randomization. In ENGAGE, for example, $T = 5$, $\left\{t_j\right\} = \{0, 4, 8, 12, 24\}$, and $t^* = t_3 = 8$. Let $\boldsymbol{X}$ be a vector of mean-centered baseline covariates, such as age at baseline, sex, etc.

We are interested in $\mathrm{E}[Y_t^{(a_1, a_{2R}, a_{2NR})} \mid \boldsymbol{X}]$, the marginal mean outcome at time t under DTR (a_1, a_{2R}, a_{2NR}) conditional on $\boldsymbol{X}$. This is the mean of Y_t had all individuals been

offered DTR (a_1, a_{2R}, a_{2NR}), conditional on $\boldsymbol{X}$. Recall that a DTR recommends treatments for both responders and non-responders; therefore, $\mathrm{E}[Y_t^{(a_1,a_{2R},a_{2NR})} \mid \boldsymbol{X}]$ is marginal over response status.

Assume $\mu_t^{(a_1,a_{2R},a_{2NR})}(\boldsymbol{X};\boldsymbol{\theta})$ is a model, indexed by unknown parameters $\boldsymbol{\theta}$, for the conditional mean $\mathrm{E}[Y_t^{(a_1,a_{2R},a_{2NR})} \mid \boldsymbol{X}]$. As discussed by Lu et al. (2016), the sequential nature of treatment delivery in SMARTs may suggest constraints on the form of $\mu_t^{(a_1,a_{2R},a_{2NR})}(\boldsymbol{X};\boldsymbol{\theta})$. The form of μ_t will depend, in part, on the design of the SMART. For instance, in ENGAGE, at time $t = 0$, no treatments have been assigned, so all DTRs share a common mean. At times $t = 4$ and $t = 8$, the four embedded DTRs differ only by recommended first-stage treatment; thus there are two means of $Y_t^{(a_1,a_{2R},a_{2NR})}$ at each timepoint. Finally, for times $t > t^* = 8$, each DTR has a different mean $Y_t^{(a_1,a_{2R},a_{2NR})}$.

Any linear model which respects the design constraints described above is permissible for $\boldsymbol{\mu}^{(a_1,a_{2R},a_{2NR})}(\boldsymbol{X};\boldsymbol{\theta})$, including, for example, flexible spline models. As an example, a simple approach to modeling data from ENGAGE would be to use a piecewise linear time trend:

$$\begin{aligned}\mu_t^{(a_1,a_{2R},a_{2NR})}(\boldsymbol{X};\boldsymbol{\theta}) = \boldsymbol{\eta}^\top \boldsymbol{X} + \gamma_0 + \mathbb{1}_{\{t\le t^*\}}(\gamma_1 t + \gamma_2 a_1 t) \\ + \mathbb{1}_{\{t>t^*\}}\big(t^*\gamma_1 + t^*\gamma_2 a_1 + \gamma_3(t-t^*) + \gamma_4(t-t^*)a_1 \\ + \gamma_5(t-t^*)a_{2NR} + \gamma_6(t-t^*)a_1 a_{2NR}\big),\end{aligned} \tag{2.1}$$

where $\boldsymbol{\theta} = (\boldsymbol{\eta}^\top, \boldsymbol{\gamma}^\top)^\top$ and $\mathbb{1}_{\{E\}}$ is the indicator function for the event E. Model (2.1) is plotted in figure 2.1.

Using contrast coding, i.e., $a_i \in \{-1, 1\}$ for $i = 1, 2$, based on model (2.1),

$$2\gamma_2 = E\left[\frac{Y_{t_j}^{(1,0,\cdot)} - Y_{t_k}^{(1,0,\cdot)}}{t_j - t_k} - \frac{Y_{t_j}^{(-1,0,\cdot)} - Y_{t_k}^{(-1,0,\cdot)}}{t_j - t_k} \mid \boldsymbol{X}\right], \quad t_j, t_k \le t^*, \tag{2.2}$$

for example, represents the difference in slopes of expected treatment readiness in the first stage of the SMART between DTRs starting with different first-stage treatments. Other possible models may include covariate-by-treatment interaction terms, potentially leading to expressions in which equation (2.2) would be a function of $\boldsymbol{X}$.

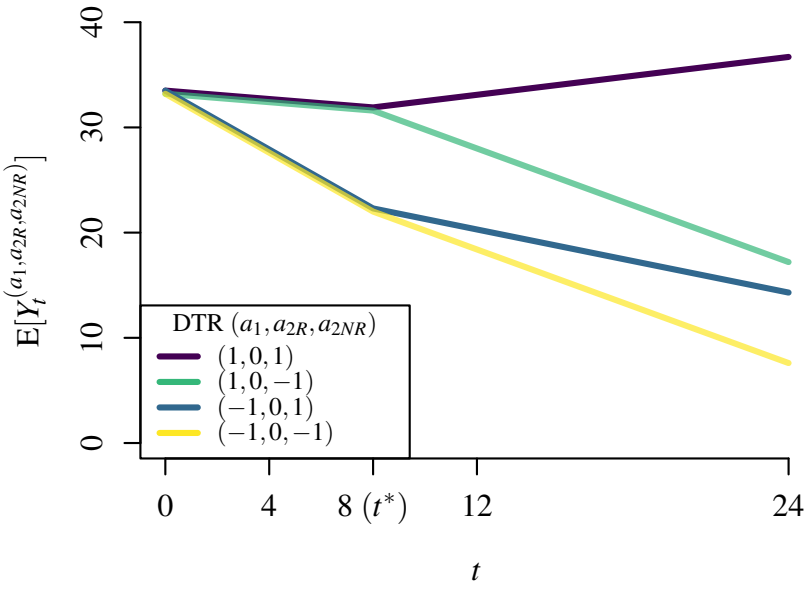

Figure 2.1: Plot of example marginal structural mean model (2.1) for a particular choice of $\boldsymbol{\gamma}$ and $\boldsymbol{\eta} = 0$. The branching structure reflects the choice of a piecewise linear model for the marginal mean.

For design I, an example marginal structural mean model is

$$\begin{aligned}\mu_t^{(a_1,a_{2R},a_{2NR})}(\boldsymbol{X};\boldsymbol{\theta}) = \boldsymbol{\eta}^\top \boldsymbol{X} + \gamma_0 + \mathbb{1}_{\{t\le t^*\}}(\gamma_1 t + \gamma_2 t a_1) \\ + \mathbb{1}_{\{t>t^*\}}\big(\gamma_1 t^* + \gamma_2 t^* a_1 + \gamma_3(t-t^*) + \gamma_4(t-t^*)a_1 + \gamma_5(t-t^*)a_{2R} \\ + \gamma_6(t-t^*)a_{2NR} + \gamma_7(t-t^*)a_1 a_{2R} + \gamma_8(t-t^*)a_1 a_{2NR}\big),\end{aligned} \tag{2.3}$$

and for design III, an example model is

$$\begin{aligned}\mu_t^{(a_1,a_{2R},a_{2NR})}(\boldsymbol{X};\boldsymbol{\theta}) = \boldsymbol{\eta}^\top \boldsymbol{X} + \gamma_0 + \mathbb{1}_{\{t\le t^*\}}(\gamma_1 t + \gamma_2 a_1 t) \\ + \mathbb{1}_{\{t>t^*\}}\big(\gamma_1 t^* + \gamma_2 t^* a_1 + \gamma_3(t-t^*) + \gamma_4(t-t^*)a_1 \\ + \gamma_5(t-t^*)\mathbb{1}_{\{a_1=1\}}a_{2NR}\big).\end{aligned} \tag{2.4}$$

2.2 Estimation

2.2.1 Observed Data

Suppose we have data arising from a SMART with n participants. Let $A_{1,i} \in \{-1,1\}$ denote the first-stage treatment randomly assigned to participant i ($i = 1,\ldots,n$), and let $R_i \in \{0,1\}$ indicate whether the ith participant responded to $A_{1,i}$, in which case $R_i = 1$, or not, so $R_i = 0$. Define $A_{2,i} \in \{-1,1\}$ to be the randomly-assigned second-stage treatment. In design II, since only non-responders are re-randomized, we set $A_{2,i} = 0$ for responders; similarly for design III. We observe a continuous outcome $Y_{t,i}$ for each participant at each of T timepoints. In general, the data collected on the ith individual over the course of the study are of the form

$$\left(\boldsymbol{X}_i, A_{1,i}, R_i, A_{2,i}, \boldsymbol{Y}_i\right),$$

where $\boldsymbol{Y}_i$ is a length-T vector consisting of all values of the outcome observed for the i^{th} participant.

2.2.2 Estimating Equations

Our goal is to estimate and make inferences on $\boldsymbol{\theta}$. Let $\mathscr{D}$ be the set of DTRs embedded in the SMART under study; for instance, in design II,

$$\mathscr{D} = \left\{(a_1, a_{2R}, a_{2NR}) : a_1 \in \{-1,1\}, a_{2R} = 0, a_{2NR} \in \{-1,1\}\right\}.$$

Let $W^{(d)}(A_{1,i}, R_i, A_{2,i})$ be a weight associated with participant i and DTR $d \in \mathscr{D}$ defined as

$$W^{(d)}(A_{1,i}, R_i, A_{2,i}) = \frac{I^{(d)}(A_{1,i}, R_i, A_{2,i})}{P(A_{1,i} = a_1)P(A_{2,i} = a_2 \mid A_{1,i} = a_1, R_i)}, \tag{2.5}$$

where $I^{(d)}(A_{1,i}, R_i, A_{2,i})$ is an indicator of whether participant i is consistent with DTR d. Note that the form of $I^{(d)}(A_{1,i}, R_i, A_{2,i})$ depends on the particular SMART design under study; for each of the designs in figure 1.1, these expressions are shown in table 2.1. $W^{(d)}(A_{1,i}, R_i, A_{2,i})$ is an inverse-probability-of-treatment weight used to correct for known imbalance in the proportion of responders and non-responders consistent with each DTR (Cole and Hernn 2008; Nahum-Shani et al. 2012a; Chakraborty and Moodie 2013). In design II, for example, only non-responders to first-stage treatment are re-randomized; if all randomizations are with probability 0.5, $W^{(1,0,1)}(1,1,0) = (.5 \times 1)^{-1} = 2$ and $W^{(1,0,1)}(1,0,1) = (.5 \times .5)^{-1} = 4$. Note that, for example, $W^{(1,0,1)}(1,0,-1) = 0$, since an individual with $A_1 = 1$, $R = 0$ and $A_2 = -1$ is not consistent with DTR $(1,0,1)$.

Table 2.1: Design-specific indicators for consistency with a given DTR $d = (a_1, a_{2R}, a_{2NR}) \in \mathscr{D}$.

Design	Form of $I^{(d)}$
I	$\mathbb{1}_{\{A_{1,i}=a_1\}}(\mathbb{1}_{\{A_{2,i}=a_{2R}\}}R_i + \mathbb{1}_{\{A_{2,i}=a_{2NR}\}}(1-R_i))$
II	$\mathbb{1}_{\{A_{1,i}=a_1\}}(R_i + \mathbb{1}_{\{A_{2,i}=a_{2NR}\}}(1-R_i))$
III	$\mathbb{1}_{\{A_{1,i}=a_1\}}\left(\mathbb{1}_{\{a_1=-1\}} + \mathbb{1}_{\{a_1=1\}}(R_i + \mathbb{1}_{\{A_{2,i}=a_{2NR}\}}(1-R_i))\right)$

Define $\boldsymbol{D}^{(d)}(\boldsymbol{X}_i)$ to be the partial derivative of $\boldsymbol{\mu}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}^\top$. Let $\boldsymbol{V}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\tau})$ be a working covariance matrix for $\boldsymbol{Y}^{(d)}$, conditional on baseline covariates $\boldsymbol{X}$, under DTR $d \in \mathscr{D}$; we discuss this quantity in detail in section 2.2.3. We estimate $\boldsymbol{\theta}$ by solving the estimating equations

$$\mathbf{0} = \frac{1}{n}\sum_{i=1}^{n}\sum_{d\in\mathscr{D}}\left[W^{(d)}\left(A_{1,i},R_i,A_{2,i}\right)\cdot\boldsymbol{D}^{(d)}(\boldsymbol{X}_i)^\top\boldsymbol{V}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\tau})^{-1}\left(\boldsymbol{Y}_i-\boldsymbol{\mu}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\theta})\right)\right] \quad (2.6)$$

for $\boldsymbol{\theta}$. We call the solution to equation (2.6) $\hat{\boldsymbol{\theta}}$.

Under usual regularity conditions for M-estimators and given data from a SMART, $\hat{\boldsymbol{\theta}}$ is asymptotically consistent for $\boldsymbol{\theta}^*$, the true regression parameter vector. Under additional assumptions described in appendix A, we can interpret $\boldsymbol{\theta}^*$ causally. Furthermore, $\sqrt{n}\left(\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}^*\right)$ has an asymptotic multivariate normal distribution:

$$\sqrt{n}\left(\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}^*\right) \Rightarrow \mathscr{N}\left(\mathbf{0},\boldsymbol{B}^{-1}\boldsymbol{M}\boldsymbol{B}^{-1}\right),$$

where

$$\boldsymbol{B} := \mathrm{E}\left[\sum_{d\in\mathscr{D}}W^{(d)}\left(A_{1,i},R_i,A_{2,i}\right)\boldsymbol{D}^{(d)}(\boldsymbol{X}_i)^\top\boldsymbol{V}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\tau})^{-1}\boldsymbol{D}^{(d)}(\boldsymbol{X}_i)\right] \quad (2.7)$$

and

$$\boldsymbol{M} := \mathrm{E}\left[\left(\sum_{d\in\mathscr{D}}W^{(d)}\left(A_{1,i},R_i,A_{2,i}\right)\boldsymbol{D}^{(d)}(\boldsymbol{X}_i)^\top\boldsymbol{V}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\tau})^{-1}\left(\boldsymbol{Y}_i-\boldsymbol{\mu}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\theta}^*)\right)\right)^{\otimes 2}\right], \quad (2.8)$$

with $\boldsymbol{Z}^{\otimes 2} = \boldsymbol{Z}\boldsymbol{Z}^\top$. Note that $\hat{\boldsymbol{\theta}}$ is consistent for $\boldsymbol{\theta}^*$ regardless of the choice of $\boldsymbol{V}^{(a_1,a_2)}(\boldsymbol{X}_i;\boldsymbol{\tau})$; however, we conjecture that choices closer to the true matrix will yield more efficient estimates.

2.2.3 Estimation of the Working Covariance Matrix

In general, for an embedded DTR $d \in \mathscr{D}$, $\boldsymbol{V}^{(d)}(\boldsymbol{X}_i; \boldsymbol{\sigma}, \boldsymbol{\rho})$ takes the form

$$\boldsymbol{V}^{(d)}(\boldsymbol{X}_i; \boldsymbol{\sigma}, \boldsymbol{\rho}) = \boldsymbol{S}^{(d)}(\boldsymbol{\sigma})^{1/2} \boldsymbol{R}^{(d)}(\boldsymbol{\rho}) \boldsymbol{S}^{(d)}(\boldsymbol{\sigma})^{1/2}, \tag{2.9}$$

where $\boldsymbol{S}^{(d)}(\boldsymbol{\sigma})^{1/2} \in \mathbb{R}^{T \times T}$ is a diagonal matrix with diagonal entries $\sigma_{t_1}^{(d)}, \ldots, \sigma_{t_T}^{(d)}$, and $\boldsymbol{R}^{(d)}(\boldsymbol{\rho})$ is a working correlation matrix for $\boldsymbol{Y}^{(d)}$. Note that this notation allows for different working covariance structures for each DTR, as well as non-constant variances in the repeated-measures outcome.

We propose the following procedure to estimate $\boldsymbol{V}^{(d)}(\boldsymbol{X}_i; \boldsymbol{\tau})$, where $\boldsymbol{\tau} = (\boldsymbol{\sigma}^\top, \boldsymbol{\rho}^\top)^\top$. First, estimate $\boldsymbol{\theta}$ by solving equation (2.6) using the $T \times T$ identity matrix as $\boldsymbol{V}^{(d)}(\boldsymbol{X}_i; \boldsymbol{\tau})$ for all $d \in \mathscr{D}$. Call the solution $\hat{\boldsymbol{\theta}}_{(0)}$. Next, use $\hat{\boldsymbol{\theta}}_{(0)}$ to estimate $\sigma_t^{(d)}$ as follows

$$\hat{\sigma}_t^{(d)} = \frac{\sum_{i=1}^n W^{(d)}\left(A_{1,i}, R_i, A_{2,i}\right) \left(Y_{i,t} - \mu_t^{(d)}\left(\boldsymbol{X}_i; \hat{\boldsymbol{\theta}}_{(0)}\right)\right)^2}{\sum_{i=1}^n W^{(d)}\left(A_{1,i}, R_i, A_{2,i}\right) - p}, \tag{2.10}$$

where p is the dimension of $\boldsymbol{\theta}$. If the scientist believes that this variance is constant over time for each DTR, the estimate in equation (2.10) can be averaged over time; one can also average over DTR if one believes the variance is constant across all embedded DTRs. Estimators for $\boldsymbol{\rho}^{(d)}$ vary with choice of correlation structure $\boldsymbol{R}^{(d)}(\boldsymbol{\rho})$; we present estimators for selected structures in table 2.2.

We estimate $\boldsymbol{V}^{(d)}(\boldsymbol{X}_i; \boldsymbol{\tau})$ by plugging appropriately-pooled estimates of $\boldsymbol{\sigma}$ from equation (2.10) and of $\boldsymbol{\rho}$ from table 2.2 into equation (2.9). The form of $\boldsymbol{R}^{(d)}(\boldsymbol{\rho})$ can be chosen according to existing domain knowledge for primary analyses; secondary analyses might use exploratory methods to discover an appropriate working correlation structure.

Table 2.2: Correlation estimators for selected working correlation structures. The top entries define estimators under the assumption of constant within-person variance over time; the bottom entries allow for time-varying variances. $d \in \mathscr{D}$ is an embedded DTR, $W_i^{(d)}$ is shorthand for $W^{(d)}(A_{1,i}, R_i, A_{2,i})$, and $\hat{e}_{i,t}^{(d)}(\hat{\boldsymbol{\theta}})$ is the estimated residual $Y_{i,t} - \hat{\mu}_t^{(d)}(\boldsymbol{X}_i; \hat{\boldsymbol{\theta}})$.

Cor. structure	$\text{Cor}(Y_{t_j}^{(d)}, Y_{t_k}^{(d)})$	Estimator
AR(1)	$\begin{cases} 1 & t_j = t_k \\ \left(\rho^{(d)}\right)^{\lvert j-k \rvert} & t_j \neq t_k \end{cases}$	$\hat{\rho}^{(d)} = \frac{\sum_{i=1}^n W_i^{(d)} \sum_{m=1}^{T-1} \hat{e}_{i,t_m}^{(d)}(\hat{\boldsymbol{\theta}}) \hat{e}_{i,t_{m+1}}^{(d)}(\hat{\boldsymbol{\theta}})}{\left(\hat{\sigma}^{(d)}\right)^2 \cdot n \cdot (T-1)}$
Exchangeable	$\begin{cases} 1 & t_j = t_k \\ \rho^{(d)} & t_j \neq t_k \end{cases}$	$\hat{\rho}^{(d)} = \frac{\sum_{i=1}^n W_i^{(d)} \sum_{l<m} \hat{e}_{i,t_l}^{(d)}(\hat{\boldsymbol{\theta}}) \hat{e}_{i,t_m}^{(d)}(\hat{\boldsymbol{\theta}})}{\left(\hat{\sigma}^{(d)}\right)^2 \cdot n \cdot T(T-1)/2}$
Unstructured	$\begin{cases} 1 & t_j = t_k \\ \rho_{t_j,t_k}^{(d)} & t_j \neq t_k \end{cases}$	$\hat{\rho}_{t_j,t_k}^{(d)} = \frac{\sum_{i=1}^n W_i^{(d)} \hat{e}_{i,t_j}^{(d)}(\hat{\boldsymbol{\theta}}) \hat{e}_{i,t_k}^{(d)}(\hat{\boldsymbol{\theta}})}{\left(\hat{\sigma}^{(d)}\right)^2 \cdot n}$

2.2.4 Iterated Estimation Procedure

After estimating $\boldsymbol{V}^{(d)}(\boldsymbol{X}_i; \boldsymbol{\tau})$, we again solve equation (2.6), this time using $\hat{\boldsymbol{V}}^{(d)}(\boldsymbol{X}_i; \hat{\boldsymbol{\tau}})$ as the working covariance matrix. This yields a "one-step" estimator of $\boldsymbol{\theta}$, which we denote by $\hat{\boldsymbol{\theta}}_{(1)}$. This process can be further iterated, as suggested by Liang and Zeger (1986); that is, we can use $\hat{\boldsymbol{\theta}}_{(1)}$ in equation (2.10) to obtain a new estimate for the working covariance structure, and so on until convergence. We call the final estimate of the model parameters $\hat{\boldsymbol{\theta}}$.

Work by Lipsitz et al. (2017) indicates that the one-step estimator is asymptotically equivalent to the "fully-iterated" estimator and is much less computationally intensive when the number of repeated measures is large. Anecdotally, we have found in reasonable simulation models for SMARTs with five or fewer measurement occasions that fully-iterated estimates tend to converge in L_2-norm within a tolerance of 10^{-8} after about six iterations and do not represent significant computational burden.

2.3 Sample Size Formulae for End-of-Study Comparisons of Embedded DTRs in Two-Stage SMARTs

The estimation procedure presented in section 2.2 is general. The marginal structural mean model $\boldsymbol{\mu}^{(d)}(\boldsymbol{X}_i; \boldsymbol{\theta})$ can take any form appropriate for the SMART under analysis, data can be observed at any number of timepoints, and the working covariance matrix can have arbitrary structure (Lu et al. 2016).

We now develop sample size formulae for SMARTs in which the primary aim is to compare the mean end-of-study outcomes for two embedded DTRs that recommend different first-stage treatments and which satisfy certain design constraints. For a variety of reasons, there is an interest in collecting repeated-measures outcomes even in settings in which the primary aim is an end-of-study comparison. Because repeated measurements within the same person are often positively correlated, analyses which leverage this information can be more efficient than those which do not (Cook and Ware 1983). This can be especially beneficial in situations with small signal-to-noise ratios. Furthermore, longitudinal data allows investigators to examine trajectories over time, regardless of the primary comparison. This can help tell a fuller story about change over time.

For the sample size methods developed here, we restrict our focus to two-stage SMARTs in which the outcome is observed at three equally-spaced measurement occasions – baseline, just prior to assessment of the tailoring variable and subsequent randomization, and at the end of the study – and in which all randomizations occur with probability 0.5. For simplicity, we ignore baseline covariates; this is a conservative assumption, since it will inflate the variance of the estimates from equation (2.6). Additionally, we consider a saturated, piecewise-linear mean structure $\boldsymbol{\mu}^{(d)}(\boldsymbol{\theta})$ similar to models (2.1), (2.3) and (2.4).

Let $\boldsymbol{c}$ be some contrast vector so that the primary aim null hypothesis takes the form

$$H_0: \quad \boldsymbol{c}^\top \boldsymbol{\theta} = \mathbf{0},$$

which we will test against an alternative of the form H_1: $\boldsymbol{c}^\top \boldsymbol{\theta} = \Delta$. To compare mean end-of-study outcomes between two embedded DTRs which recommend different first-stage treatments, the estimand of interest is

$$\boldsymbol{c}^\top \boldsymbol{\theta} = \mathrm{E}\left[Y_2^{(1,a_{2\mathrm{R}},a_{2\mathrm{NR}})} - Y_2^{(-1,a'_{2\mathrm{R}},a'_{2\mathrm{NR}})} \right], \tag{2.11}$$

for some choice of $a_{2\mathrm{R}}$, $a'_{2\mathrm{R}}$, $a_{2\mathrm{NR}}$, and $a'_{2\mathrm{NR}}$. For example, to test equality of mean end-of-study outcomes for DTRs (1, 0, 1) and (-1, 0, -1) in design II under model (2.1) (assuming no $\boldsymbol{X}$, $\{t_j\} = \{0,1,2\}$, $t^* = 1$), the estimand (2.11) can be written as the linear combination $\boldsymbol{c}^\top \boldsymbol{\gamma}$, where $\boldsymbol{c}^\top = (0,0,2,0,2,2,0)$.

Because we are interested in a single contrast (i.e., $\boldsymbol{c}$ is a vector, not a matrix) we employ a 1-degree of freedom Wald test. The test statistic is

$$Z = \frac{\sqrt{n}\boldsymbol{c}^\top \hat{\boldsymbol{\theta}}}{\sigma_c},$$

where $\sigma_c = \sqrt{\boldsymbol{c}^\top \boldsymbol{B}^{-1}\boldsymbol{M}\boldsymbol{B}^{-1}\boldsymbol{c}}$. Under the null hypothesis, by asymptotic normality of $\sqrt{n}\left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\right)$, the test statistic follows a standard normal distribution.

Define δ to be the standardized effect size as described by Cohen (1988) for an end-of-study comparison, i.e.,

$$\delta = \frac{\Delta}{\sigma}, \tag{2.12}$$

where $\sigma = \text{Var}(Y_t^{(d)})$ (see working assumption A2 below).

The sample size formulae will require the response rate $P(R^{(a_1)} = 1) = r_{a_1}$. In order to simplify the form of σ_c and obtain tractable, intepretable sample size formulae, we make the following working assumptions:

A1 *Constrained conditional covariance matrices for DTRs under comparison.*

(a) The variability of $Y_t^{(d)}$ around the DTR mean $\mu_t^{(d)}(\boldsymbol{\theta})$ among responders and non-responders is no more than the variance of $Y_t^{(d)}$ unconditional on response, i.e.,

$$\text{E}\left[\left(Y_t^{(d)} - \mu_t^{(d)}(\boldsymbol{\theta})\right)^2 \mid R^{(a_1)}\right] \leq \text{E}\left[\left(Y_t^{(d)} - \mu_t^{(d)}(\boldsymbol{\theta})\right)^2\right],$$

for all $t > t^*$ and DTRs $d \in \mathscr{D}$ under study.

(b) For times $t_i \leq t_j \leq t^*$, response status is uncorrelated with products of residuals of Y_{t_i}, i.e,

$$\text{Cov}\left(R^{(a_1)}, \left(Y_{t_i}^{(d)} - \mu_{t_i}^{(d)}(\boldsymbol{\theta})\right)\left(Y_{t_j}^{(d)} - \mu_{t_j}^{(d)}(\boldsymbol{\theta})\right)\right) = 0.$$

for DTRs $d \in \mathscr{D}$ under study.

(c) The covariance between the end-of-study measurement and the measurements prior to the second stage among responders is less than or equal to the same quantity among non-responders:

$$\text{Cov}\left(Y_t^{(d)}, Y_2^{(d)} \mid R^{(a_1)} = 1\right) \leq \text{Cov}\left(Y_t^{(d)}, Y_2^{(d)} \mid R^{(a_1)} = 0\right)$$

for DTRs $d \in \mathscr{D}$ under study and $t = 0, 1$. An additional, related assumption is given in appendix B.2.

A2 *Exchangeable marginal covariance structure.* The marginal variance of $\boldsymbol{Y}^{(d)}$ is constant across time and DTR, and has an exchangeable correlation structure with corre-

lation ρ, i.e.,

$$\text{Var}\left(\boldsymbol{Y}^{(d)}\right) = \boldsymbol{\Sigma} = \sigma^2 \begin{bmatrix} 1 & \rho & \rho \\ \rho & 1 & \rho \\ \rho & \rho & 1 \end{bmatrix}$$

for all $d \in \mathscr{D}$.

These working assumptions may be seen as overly simplifying; however, we will see in sections 2.4 and 2.5 that formula (2.13) is robust to moderate violations of working assumption A1 and that inputs to the formula can be adjusted in a way to accommodate violations of working assumption A2. A working assumption similar to A1(a) is commonly made in developing sample-size formulae for SMARTs using end-of-study outcomes (Oetting et al. 2011; Kidwell et al. 2018; NeCamp, Kilbourne, and Almirall 2017). Working assumptions A1(b) and A1(c), as well as A2, are necessary for the extension to the setting of a repeated-measures outcome with our proposed estimator.

Working assumption A1 arises specifically as a consequence of unequal weights in equation (2.6) (i.e., when there exists imbalance between responders and non-responders, by design); therefore, the assumption is not necessary in design I, and can be relaxed to apply to only the two DTRs in which non-responders are re-randomized in design III. Furthermore, working assumption A2 cannot be satisfied in design I if all eight embedded DTRs have unique means.

Under working assumptions A1 and A2, the minimum-required sample size to detect a standardized effect size δ with power at least $1-\beta$ and two-sided type-I error α is

$$n \geq \frac{4\left(z_{1-\alpha/2} + z_{1-\beta}\right)^2}{\delta^2} \cdot (1-\rho^2) \cdot \text{DE}, \tag{2.13}$$

where DE is a SMART-specific "design effect" for an end-of-study comparison (see table 2.3). Note that the first term in formula (2.13) is the typical sample size formula for a traditional two-arm randomized trial with a continuous end-of-study outcome and equal randomization probability. The middle term is due to the within-person correlation in the outcome, and is identical to the corresponding correction term for GEE analyses sized to detect a group-by-time interaction when there is no baseline group effect (see, e.g., Fitzmaurice, Laird, and Ware (2011), ch. 20).

The sample size formula presented in formula (2.13) is conservative. It becomes more conservative as ρ approaches $(1+\sqrt{5})/2 \approx 0.62$. A sharper formula is available in appendix B.2; however, we emphasize formula (2.13) as it is more immediately intepretable.

Table 2.3: Design effects for the sample size formulae in formula (2.13). r_{a_1} is the response rate to first-stage treatment a_1

Design	Design effect	Conservative design effect
I	2	2
II	$\frac{1}{2}(2-r_1)+\frac{1}{2}(2-r_{-1})$	2
III	$\frac{1}{2}(3-r_1)$	$\frac{3}{2}$

2.4 Simulation Results

Initial simulation results are compiled in table 2.4. Due to recent updates to the working assumptions in section 2.3 in order to make them more realistic, the data-generative model on which table 2.4 is based does not allow for violations of working assumption A1(b). More details on this can be found in section 3.1.1.

We find that sample size formula (2.13) performs as expected when all assumptions are satisfied. Empirical power is not significantly less than the target power of 0.8, per a one-sided binomial test with level 0.05. The sample size is, as expected, slightly conservative, particularly when within-person correlation is high. There may be some concern that, for high within-person correlation, formula (2.13) is overly conservative; should this concern arise, we recommend use of the sharper formula presented in appendix B.2.

Violation of working assumption A1(a) was induced by lowering the end-of-study variance among responders relative to that among non-responders, while keeping the marginal variance fixed. In particular, the results shown in table 2.4 correspond to approximately a 25% reduction in responders' variance relative to the non-responders' variance for all DTRs.

As conjectured in section 2.3, violating working assumption A1(a) does not impact empirical power in design I, since the assumption arises as a consequence of imbalanced numbers of responders and non-responders consistent with a particular DTR. For design II, empirical power is consistently less than the nominal value when working assumption A1(a) is violated. However, while the empirical power is often (statistically) significantly less than 0.8, the observed loss of power is relatively small. For design III, we notice small reductions in power relative to scenarios in which both working assumptions A1 and A2 are satisfied, though the conservative nature of formula (2.13) appears to protect against more severe loss of power. This suggests that our sample size formula is moderately robust to "reasonable" violations of A1(a).

Violation of working assumption A1(c) was induced by choosing $\text{Cor}(Y_t^{(d)}, Y_2^{(d)} \mid R^{(a_1)} = 1) > \text{Cor}(Y_t^{(d)}, Y_2^{(d)} \mid R = 0)$ while keeping respective variances fixed. There exist natural

constraints on how much larger than $\text{Cov}(Y_t^{(d)}, Y_2^{(d)} \mid R = 0)$ the responders' covariance can be while ensuring that (1) all conditional covariance matrices are positive definite and (2) $\text{Cov}(Y_t^{(d)}, Y_2^{(d)} \mid R = 0) \geq 0$ for $t = 0, 1$. These constraints vary with ρ. The empirical power results shown in table 2.4 were generated by choosing $\text{Cor}(Y_t^{(d)}, Y_2^{(d)} \mid R^{(a_1)} = 1)$ as the midpoint between the minimum covariance for which the assumption is violated and the maximum covariance allowed by the aforementioned constraints. Simulation results show that our sample size formula is quite robust to violations for low-to-moderate within-person correlations; at high correlations, the empirical power is significantly less than 0.8. However, as with working assumption A1(a), the observed reduction in power is not unreasonable. Furthermore, when within-person correlation is high, sample size becomes rather small. Since the method presented here is based on asymptotic normality, we caution the reader that small sample sizes (e.g., $n < 100$) provided by formula (2.13) may be quite sensitive to violation of the working assumptions.

The final columns of table 2.4 suggest that formula (2.13) is highly sensitive to violations of working assumption A2 in regards to the true correlation structure. In particular, when the true correlation structure is not exchangeable with correlation ρ and is instead AR(1) with correlation ρ, empirical power is substantially lower than the target of 0.8, particularly as ρ increases. This is unsurprising: under an AR(1) correlation structure, less information about the end-of-study outcome is provided by, say, the baseline measure than under an exchangeable correlation structure. Since, by using formula (2.13), we have assumed more information is available from earlier measurements than is actually the case, we will be underpowered. As our assumed ρ increases, the difference between the assumed and actual correlation between the end-of-study measurement and earlier measurements increases, leading to more severe loss of power.

In figure 2.2, we examine the effect on empirical power of misspecifying the within-person correlation. Analytically, we see from formula (2.13) that if the assumed ρ is smaller than the true within-person correlation, the sample size will be conservative. On the other hand, when the assumed ρ in formula (2.13) is larger than the true correlation, the sample size will be anti-conservative. Figure 2.2 shows plots of empirical power against the difference between the assumed within-person correlation ρ_{guess} and the true ρ. For small ρ_{guess}, formula (2.13) appears to be quite robust to misspecification of ρ; however, as ρ_{guess} increases, the formula becomes highly sensitive to such a violation of working assumption A2. This is supported analytically, since formula (2.13) is a function of ρ_{guess}^2.

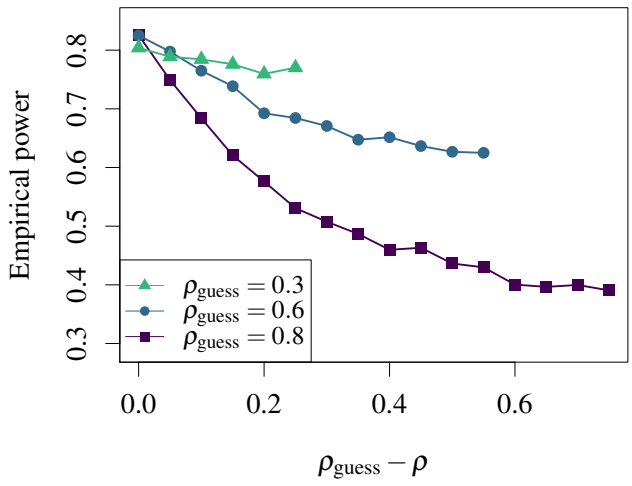

Figure 2.2: Empirical power versus the difference between the true within-person correlation ρ and hypothesized correlation ρ_{guess} used to compute sample size. Results are shown for design II with a hypothesized response rate of 0.4, and sample size was chosen to detect standardized effect size $\delta = 0.3$ for the comparison of DTRs (1, 0, 1) and (-1, 0, -1). Each point is based on 5000 simulations with target power 0.8 and significance level 0.05. Results are similar for designs I and III and different values of δ and r.

Table 2.4: Sample sizes and empirical power results for an end-of-study comparison of the DTR recommending only treatments indexed by 1 and that which recommends only treatments indicated by -1. δ is the standardized effect size defined in equation (2.12). Bolded results are significantly less than 0.8 at the 0.05 level.

					Empirical power			
						Violation of A1		Violation of A2
Design	δ	r	ρ	n	A1 and A2 satisfied	A1(a)	A1(c)	True AR(1)
I	0.3	0.4	0	698	0.797	0.800	–	–
			0.3	635	0.807	0.811	0.820	**0.778**
			0.6	447	0.842	0.829	0.830	**0.712**
			0.8	252	0.848	0.838	0.844	**0.662**
		0.6	0	698	0.816	0.794	–	–
			0.3	635	0.825	0.801	0.813	**0.778**
			0.6	447	0.829	0.833	0.833	**0.723**
			0.8	252	0.851	0.832	0.838	**0.665**
	0.5	0.4	0	252	0.804	0.810	–	–
			0.3	229	0.818	0.812	0.820	**0.783**
			0.6	161	0.843	0.829	0.837	**0.710**
			0.8	91	0.845*	0.840*	0.840*	**0.676***
		0.6	0	252	0.809	0.797	–	–
			0.3	229	0.816	0.818	0.812	**0.777**
			0.6	161	0.838	0.831	0.831	**0.713**
			0.8	91	0.853	0.840	0.846	**0.666**
II	0.3	0.4	0	559	0.800	**0.790**	–	–
			0.3	508	0.803	**0.786**	**0.785**	**0.757**
			0.6	358	0.824	0.795	**0.779**	**0.695**
			0.8	201	0.825	**0.785**	0.803	**0.625**
		0.6	0	489	0.796	**0.773**	–	–
			0.3	445	0.797	**0.787**	**0.786**	**0.767**
			0.6	313	0.812	**0.783**	**0.766**	**0.679**
			0.8	176	0.827	**0.756**	**0.774**	**0.625**
	0.5	0.4	0	201	0.794	0.793	–	–
			0.3	183	0.815	0.794	**0.789**	**0.774**
			0.6	129	0.830	0.797	0.793	**0.699**
			0.8	73	0.839	**0.787**	0.807	**0.638**
		0.6	0	176	0.806	**0.765**	–	–
			0.3	160	0.815	**0.773**	0.802	**0.778**
			0.6	113	0.816*	**0.773**	**0.763**	**0.691**
			0.8	64	0.831*	**0.775**	**0.787**	**0.643**
III	0.3	0.4	0	454	0.798	0.793	–	–
			0.3	413	0.805	0.800	0.800	**0.760**
			0.6	291	0.808	0.803	0.797	**0.677**
			0.8	164	0.825	0.800	0.802	**0.611**
		0.6	0	419	0.798	0.805	–	–
			0.3	381	0.802	0.793	0.795	**0.753**
			0.6	268	0.814	0.803	**0.786**	**0.686**
			0.8	151	0.824	0.794	**0.784**	**0.611**
	0.5	0.4	0	164	0.802	**0.790**	–	–
			0.3	149	0.814	0.803	0.805	**0.773**
			0.6	105	0.815	0.807	0.796	**0.683**
			0.8	59	0.811	0.815*	0.817*	**0.635***
		0.6	0	151	0.792	0.791	–	–
			0.3	138	0.813	0.802	0.799	**0.769**
			0.6	97	0.818*	0.804*	0.796*	**0.690***
			0.8	55	0.824*	0.797*	0.797*	**0.630***

* Fewer than 5000 simulations generated data in which all treatment sequences were observed.

2.5 Discussion

We have derived sample size formulae for SMART designs in which the primary aim is a comparsion of two embedded DTRs that begin with different first-stage treatments on a continuous, repeated-measures outcome. We derived the formulae for three common SMART designs.

The sample size formula is the product of three components: (1) the formula for the minimum sample size for the comparison of two means in a standard two-arm trial (see, e.g., Friedman, Furberg, and DeMets (2010) page 147), (2) a deflation factor of $1-\rho^2$ that accounts for the use of a repeated-measures outcome, and (3) a SMART-specific "design effect", an inflation factor that accounts for the SMART design.

The SMART design effect can be interpreted as the cost of conducting the SMART relative to conducting a standard two-arm randomized trial of the two DTRs which comprise the primary aim. The benefit of conducting a SMART (relative to the standard two-arm randomized trial) is the ability to answer additional, secondary questions that are useful for constructing effective DTRs (Almirall et al. 2018). For example, such questions may focus on one or more of the other pairwise comparisons between DTRs, on whether the first- and second-stage treatments work synergistically to impact outcomes (e.g., a test of the null that $\gamma_6 = 0$ in model (2.1)), or may focus on hypothesis-generating analyses that seek to estimate more deeply-tailored DTRs (Watkins 1989; Nahum-Shani et al. 2012b; Zhang et al. 2015).

The formulae are expected to be easy-to-use for both analysts and clinicians. Indeed, inputs α, β, and Δ are as in the sample size formula for a standard z-test. Furthermore, estimates of ρ, r_{a_1}, and σ are often readily available from the literature or can be estimated using data from prior studies (e.g., prior randomized trials, or external pilot studies).

We make a number of recommendations concerning the use of the formulae; in particular, how best to use the formulae conservatively in the absence of certainty concerning prior estimates of ρ, r_{a_1}, and/or the structure of the variance of the repeated measures outcome. First, in designs II and III, if there is uncertainty concerning the response rate (e.g., response rate estimates are based on data from smaller prior studies), one approach is to err conservatively by assuming a smaller-than-estimated response rate. In both designs, the most conservative approach is to assume a response rate of zero.

Second, as in standard randomized trials in which the primary aim is a pre-post comparison, the required sample size decreases as the hypothesized within-person correlation increases (Zhang, Cao, and Ahn 2014). Therefore, if the hypothesized ρ is larger than the true ρ, the computed sample size will be anti-conservative, resulting in an under-powered

study. Indeed, we see this in the results of the simulation experiment (see figure 2.2). Here, again, one approach is to err conservatively towards smaller values of ρ.

Finally, working assumption A2 (concerning the variance of the repeated measures outcome) may be seen as overly restrictive in the imposition of an exchangeable correlation structure. For example, studies with a continuous repeated measures outcome may observe an autoregressive correlation structure. However, the exchangeable working assumption can be employed conservatively in the following way: If the hypothesized structure is not exchangeable, one approach is to set ρ in formula (2.13) to the smallest plausible value. Because this approach utilizes a lower bound on the value of the true within-person correlations, it is expected to yield a larger than needed (more conservative) sample size. Similarly, if the true within-person correlation is expected to differ by DTR, one approach is to employ the smallest plausible ρ. As with the third recommendation, these recommendations are not unique to SMARTs; indeed, these strategies may also be used to size standard two-arm randomized trials with repeated measures outcome.

CHAPTER 3

Future Work

In this chapter, we propose work that will comprise the remainder of the dissertation. We begin by describing what remains to be done before publication of chapter 2, as well as immediate follow-up work involving software development. Next, we propose two extensions to chapter 2. In the first, we propose to investigate a trade-off between sample size and the number of measurement occasions in SMARTs with repeated-measures outcomes. In the second, we discuss extending our longitudinal modeling framework to accommodate intensive longitudinal data using a flexible spline approach.

3.1 Next Steps for Sample Size Calculation in Repeated-Measures SMARTs

3.1.1 Work Remaining Prior to Manuscript Submission

As noted in section 2.4, recent revisions to make working assumptions A1 and A2 more realistic have necessitated a redeveloped generative model for simulations. While it is straightforward to design simulations in which working assumptions A1 and A2 are satisfied, it is more difficult to build a data-generative model in which the assumptions are violated in a controlled way. The results tabulated in table 2.4 are from a generative model that assumes response is independent of Y_0 and Y_1. While this clearly satisfies working assumption A1(b), which is of interest when assessing the robustness of our formulae to other assumptions, it remains to discover how to violate working assumption A1(b) to a known extent without violating other assumptions.

In the case where $\text{Var}(Y_t^{(d)})$ varies with time and/or DTR, we conjecture that empirical power will be less than the nominal value if a pooled estimate of σ^2 is used when the variance decreases with time. To see this, consider that the standardized effect size δ defined in equation (2.12) has as a denominator the standard deviation of $Y_2^{(d)}$ pooled across the groups under comparison. Should the estimate of pooled standard deviation be

larger than the true value, the variance of $\boldsymbol{c}^\top \hat{\boldsymbol{\theta}}$ will increase; since the estimate will be less efficient than hypothesized, power will be lower than expected. Conversely, we also conjecture that when $\text{Var}(Y_t^{(d)})$ increases with t, the sample size will be conservative using similar reasoning. Early simulations suggest that this conjecture may be incorrect, despite it being intuitive. This may be due to problems with a previous simulation model, and will be further investigated prior to submission.

3.1.2 Software Tools for SMARTs with Continuous Longitudinal Outcomes

Building a simulation model for this paper has resulted in extensive work on how longitudinal data from SMARTs can be generated. This includes code for both analysis and simulation. The simulation model is, by necessity, more general and flexible than those needed by other papers involving longitudinal outcomes in SMARTs, due primarily to the need to control the variance of $\boldsymbol{Y}$. This data-generative model, then, may prove to be a useful standalone tool.

While standard GEE software can be used to fit models described in section 2.1 using a so-called "weighted-and-replicated" approach (see, e.g., Nahum-Shani et al. (2012a) or the supplement to Lu et al. (2016)), it can be challenging to provide such software an appropriate working covariance model. In particular, SAS requires multiple runs of `PROC GENMOD` and `GEEPACK` in R needs a complicated user-defined working correlation structure that may be difficult to understand. We have implemented in R a wrapper to a standard Newton-Raphson solver designed to fit models such as those described in section 2.1. The wrapper automatically estimates working variance parameters using equation (2.10) and table 2.2 and iterates the estimation procedure until convergence up to a specified tolerance. Because this solver was written with SMARTs in mind, it can be more computationally efficient than general-purpose tools while also offering a more intuitive interface for analyzing SMART data. We propose to develop an R package containing both simulation and analysis tools to make it easier to work with SMART data.

Finally, despite the relative ease of implementation of the sample size formula (2.13), development of an easy-to-use, flexible software tool would extend the practical impact of this work. This would likely take the form of a web-based sample size calculator developed in Shiny (Chang et al. 2018). This would allow practitioners to, for instance, specify non-exchangeable working correlation structures (rather than using formula (2.13) in a conservative way), and, potentially, incorporate additional timepoints (see section 3.3 below) and other summary measures of the longitudinal outcome (e.g., area under the

curve). The interface would be modeled on previous sample size calculators we have developed, in particular an application supporting the work of Kidwell et al. (2018) at `https://methodologycenter.shinyapps.io/SMARTsize`.

3.2 Sample Size for Repeated-Measures SMARTs with Budget Constraints

A natural extension to the work presented in chapter 2 is as follows. Consider a SMART with a repeated-measures outcome and a fixed monetary budget. Intuitively, and as we have argued in section 2.3, if the scientist believes these measurements are correlated, we expect a decrease in required sample size as the number of measurement occasions increases. This leads to a natural question: for a fixed budget, how does one choose the sample size and number of measurement occasions to maximize power? There is a broad literature exploring the selection of sample size and number of measurement occasions. To our knowledge, however, this problem has not been addressed in the context of a SMART. An extension to this setting will require consideration of weights, the fact that individuals can be consistent with more than one embedded DTR, the previously-described modeling constraints, and the requirement that there be at least one measurement occasion per stage.

Overall and Doyle (1994) provided a variety of sample size formulae for "[ANOVA] tests of significance in a two-group repeated measurements design", considering a variety of possible comparisons with and without adjustment for baseline covariates. They discovered that, for the these comparisons, the total number of repeated measurements is less meaningful for power than the number of individuals enrolled in the study. In particular, at least for ANOVA tests, "the analysis of simple difference scores at endpoint is blind to the number of intervening repeated measurements" (Overall and Doyle 1994).

Arndt et al. (2000) explored this problem in terms of "precision"; specifically, the relative contributions of additional sample size versus more measurement occasions to the standard error of the estimate for mean change over time. The work is largely empirical, confirming and offering a slight extension to that of Overall and Doyle (1994) and Maxwell (1998). The idea of optimizing against the standard error of estimates is interesting; although, for a fixed effect size, target power, and type-I error, this can be equivalent to optimizing for power.

Raudenbush and Xiao-Feng (2001) considered this problem in two-level hierarchical models, representing power for a detecting a treatment effect as a function of study duration, measurement frequency, and sample size. The authors focused primarily on "group

differences in polynomial change", modeling individual trajectories as polynomial functions in time. Furthermore, they conceptualized the effect size as a standardized mean difference in polynomial trend. This may be a useful framework to consider, as purely linear trends may be overly restrictive in some settings in which there are more than three timepoints.

Zhang and Ahn (2011a) extended this work in the context of a test for a group-by-time interaction using a GEE estimator under a cost constraint. Their results are useful for two-arm randomized trials in which the primary aim is a comparison of slopes. Though similar, the setting considered in section 2.3 is more complicated, primarily because it is difficult to derive a closed form of σ_c for a general number measurement occasions. Further work by the same authors extended the results of Overall and Doyle (1994) by considering a regression model for time-averaged outcomes across groups (Zhang and Ahn 2011b).

Future work along these lines will extend the sample size methods of chapter 2 to more than three measurement occasions, spaced throughout both stages of the trial. If a formula such as (2.13) can be achieved, we propose to discover how to optimize the selection of sample size and the number of measurement occasions, subject to a budget constraint, so as to maximize power for the SMART's primary comparison. We believe this will be an impactful contribution for trialists and practitioners designing SMARTs with repeated-measures outcomes.

3.3 SMARTs with Intensive Longitudinal Outcomes

The advent of new technology such as mobile phones and wearable sensors has contributed to an explosion of data which can be collected from an individual over a period of time. Fitness trackers can monitor and collect data on an individual's heart rate and step count nearly every minute, and smartphones provide both a source of data and a new surface on which interventions can be delivered. This is one example of the collection of intensive longitudinal data (ILD) (Walls and Schafer 2006).

It is tempting to believe that intensive longitudinal data is differentiated from traditional repeated-measures data primarily in terms of the amount of data collected. However, as noted by Walls and Schafer (2006),

> the features that make ILD unique and worthy of special consideration pertain to the scientific motivations for collecting them, the unusual nature of the hypotheses they are intended to address, and the complex features of the data that need to be revealed and parameterized by statistical models.

In contrast to more traditional repeated-measures data, ILD is observed (potentially much) more frequently and can provide more detailed information on an individual's trajectory over time. This allows researchers to study the dynamics of behavioral or disease processes on a much finer scale compared to a more conventional setting in which relatively few observations are made (Hamaker and Wichers 2017).

Because of the SMART's usefulness in constructing DTRs, which are decision rules leading to a sequence of treatments tailored to an individual's changing needs *over time*, there is increasing interest in collecting intensive longitudinal data throughout a SMART. This would enable more detailed assessment of the impact of treatment over time, as well as any delayed effects of treatment that may arise as a consequence of the sequencing that takes place within a DTR. In some settings, these effects may be quite proximal; intensive longitudinal data can potentially capture brief changes. An example of the use of ILD in SMARTs is given in section 5.3 of Lu et al. (2016), in which the authors model the outcome using regression splines. This work, particularly the details of how to incorporate design features of a SMART into the model, is discussed only briefly.

We propose to explore the use of intensive longitudinal outcomes in SMARTs, initially by formalizing guidance on how to account for the sequential nature of treatment delivery in the context of a flexible spline-based model. More traditional models with linear or polynomial time trends may prove inadequate for describing trajectories across a relatively large number of measurement occasions. However, more flexible models may require modifications of the estimation framework developed in section 2.2: while it is possible to use GEE to fit models for intensive longitudinal data (see, e.g., chapter 2 of Walls and Schafer (2006)), the sample size requirements for GEE's asymptotic properties to hold could be difficult to achieve in a study involving potentially highly burdensome data collection. Thus, there exist open questions about modeling, analytic and sample size considerations for SMARTs with a relatively large number of measurement occasions.

APPENDIX A

Identifiability Assumptions

We make the following assumptions in order to show that equation (2.6) has mean zero.

I1 *Positivity.* The probabilities $P(A_1 = 1)$ and $P(A_2 = 1 \mid A_1, R)$ are non-zero.

I2 *Consistency* (Robins 1997). A participant's observed responder status is consistent with the participant's corresponding potential responder status under the assigned first-stage treatment; i.e.,

$$R_i = \mathbb{1}_{\{A_{1,i}=1\}} R^{(1)} + \mathbb{1}_{\{A_{1,i}=-1\}} R^{(-1)}.$$

And a participant's observed repeated measures outcomes are consistent with the participant's corresponding potential repeated measures outcomes under the assigned treatment sequence; see table A.1.

I3 *Sequential randomization.* At each stage in the SMART, observed treatments A_1 and A_2 are assigned independently of future potential outcomes, given the participant's history up to that point. That is,

$$\{\boldsymbol{Y}^{(d)}_{t\leq t^*}, R(a_1)\} \perp\!\!\!\perp A_1$$
$$\{\boldsymbol{Y}^{(d)}_{t>t^*}\} \perp\!\!\!\perp A_2 \mid A_1, R$$

Identifiability assumptions I1 and I3 are satisfied by design in a SMART (see, e.g., Lavori and Dawson (2014)); identifiability assumption I2 is connects the potential outcomes and observed data, and is typically accepted in the analysis of randomized trials.

Table A.1: Design-specific consistency assumptions. $d \in \mathscr{D}$ indexes embedded DTRs (a_1, a_{2R}, a_{2NR}).

Design	Time t	$Y_{t,i}$ equals
I	t_0	$Y_{t,i}^{(d)}$
	$t_0 < t \leq t^*$	$\sum_{d\in\mathscr{D}} \mathbb{1}_{\{A_{1,i}=a_1\}} Y_{t,i}^{(d)}$
	$t > t^*$	$\sum_{d\in\mathscr{D}} \frac{1}{2}\mathbb{1}_{\{A_{1,i}=a_1\}} \left(R_i \mathbb{1}_{\{A_{2,i}=a_{2R}\}} + (1-R_i)\mathbb{1}_{\{A_{2,i}=a_{2NR}\}}\right) Y_{t,i}^{(d)}$
II	t_0	$Y_{t,i}^{(d)}$
	$t_0 < t \leq t^*$	$\sum_{d\in\mathscr{D}} \mathbb{1}_{\{A_{1,i}=a_1\}} Y_{t,i}^{(d)}$
	$t > t^*$	$\sum_{d\in\mathscr{D}} \mathbb{1}_{\{A_{1,i}=a_1\}} \left(\frac{1}{2}R_i + (1-R_i)\mathbb{1}_{\{A_{2,i}=a_2\}}\right) Y_{t,i}^{(d)}$
III	t_0	$Y_{t,i}^{(d)}$
	$t_0 < t \leq t^*$	$\sum_{d\in\mathscr{D}} \mathbb{1}_{\{A_{1,i}=a_1\}} Y_{t,i}^{(d)}$
	$t > t^*$	$\sum_{d\in\mathscr{D}} \mathbb{1}_{\{A_{1,i}=a_1\}} \left(\mathbb{1}_{\{a_1=-1\}} + \mathbb{1}_{\{a_1=1\}} \left(\frac{1}{2}R_i + (1-R_i)\mathbb{1}_{\{A_{2,i}=a_2\}}\right)\right) Y_{t,i}^{(d)}$

The factor of $1/2$ applied to some (or all) participants when $t > t*$ accounts for the fact that these participants are consistent with two DTRs. In design I, all participants are consistent with two DTRs. In design II, only responders are consistent with two DTRs, so, if $R_i = 1$ for some i, $Y_{t>t^*,i}^{(a_1,0,1)} = Y_{t>t^*,i}^{(a_1,0,-1)} := Y_{t>t^*,i}^{(a_1,0,0)}$. Similarly for responders to $a_1 = 1$ in design III.

APPENDIX B

Proofs and Derivations

B.1 Asymptotic Results

We show that $\hat{\boldsymbol{\theta}}$, the solution to equation (2.6) over $\boldsymbol{\theta}$, is asymptotically consistent for $\boldsymbol{\theta}^*$, the true regression parameter.

Define $\hat{\boldsymbol{\theta}}_n$ to be the solution of the estimating equations

$$\mathbf{0} = \frac{1}{n}\sum_{i=1}^{n}\sum_{d\in\mathscr{D}}\left[W^{(d)}\left(A_{1,i},R_i,A_{2,i}\right)\cdot \boldsymbol{D}^{(d)}(\boldsymbol{X}_i)^\top \boldsymbol{V}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\tau})^{-1}\left(\boldsymbol{Y}_i - \boldsymbol{\mu}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\theta})\right)\right] \tag{(2.6) revisited}$$

using data from n individuals. Let $\boldsymbol{Z}_i$ contain the ith individual's observed covariates (including outcome, treatment assignments, etc.). We can re-write equation (2.6) as

$$\mathbf{0} = \Psi_n(\boldsymbol{\theta}) = \frac{1}{n}\sum_{i=1}^{n}\psi_{\boldsymbol{\theta}}(\boldsymbol{Z}_i), \tag{B.1}$$

where

$$\psi_{\boldsymbol{\theta}}(\boldsymbol{Z}_i) = \sum_{d\in\mathscr{D}} W^{(d)}\left(A_{1,i},R_i,A_{2,i}\right)\cdot \boldsymbol{D}^{(d)}(\boldsymbol{X}_i)^\top \boldsymbol{V}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\tau})^{-1}\left(\boldsymbol{Y}_i - \boldsymbol{\mu}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\theta})\right).$$

Let $\hat{\boldsymbol{\theta}}_n$ be a solution to equation (B.1) for given n, and define $\boldsymbol{\theta}^*$ as the true parameter value, such that $\boldsymbol{\theta}^*$ is a zero of $\Psi(\boldsymbol{\theta}) = \mathrm{E}[\psi_{\boldsymbol{\theta}}(\boldsymbol{Z})]$.

Assuming the parameter space $\boldsymbol{\Theta}$ is compact, $\sup_{\boldsymbol{\theta}\in\boldsymbol{\Theta}}\|\Psi_n(\boldsymbol{\theta}) - \Psi(\boldsymbol{\theta})\| \xrightarrow{p} 0$ by the weak law of large numbers for random functions. If the model $\boldsymbol{\mu}^{(d)}(\boldsymbol{X}_i;\boldsymbol{\theta})$ is correctly specified and $\boldsymbol{\theta}^*$ is the unique solution of $\Psi(\boldsymbol{\theta}) = \mathrm{E}[\psi_{\boldsymbol{\theta}}(\boldsymbol{Z})]$, then consistency follows from standard results for M-estimation of a location parameter (see, e.g., Keener (2010) Theorems 9.2, 9.4, and 9.33).

We begin establishing asymptotic normality by considering a first-order Taylor expansion of the estimating equations (2.6) about $\boldsymbol{\theta}^*$, assuming continuous differentiability of

$\psi_{\boldsymbol{\theta}}$:

$$\mathbf{0} = \Psi_n(\hat{\boldsymbol{\theta}}_n) = \Psi_n\left(\boldsymbol{\theta}^*\right) + \Psi_n'\left(\tilde{\boldsymbol{\theta}}\right)\left(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\right), \tag{B.2}$$

where $\tilde{\boldsymbol{\theta}}$ is some intermediate value between $\hat{\boldsymbol{\theta}}_n$ and $\boldsymbol{\theta}^*$. Note that $\Psi_n'\left(\tilde{\boldsymbol{\theta}}\right)$ is a $p \times p$ matrix, where p is the dimension of $\boldsymbol{\theta}$. If $\Psi_n'\left(\tilde{\boldsymbol{\theta}}\right)$ is non-singular, equation (B.2) can be re-written as

$$\sqrt{n}\left(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\right) = -\sqrt{n}\Psi_n'\left(\tilde{\boldsymbol{\theta}}\right)^{-1}\Psi_n\left(\boldsymbol{\theta}^*\right). \tag{B.3}$$

By the central limit theorem, $\sqrt{n}\Psi_n\left(\boldsymbol{\theta}^*\right) \Rightarrow \mathcal{N}\left(\mathbf{0}, \mathrm{Var}(\psi_{\boldsymbol{\theta}}(\boldsymbol{Z}))\right)$.

Under sufficient regularity conditions (see, e.g., van der Vaart (1998) theorem 5.41), and because $\tilde{\boldsymbol{\theta}} \xrightarrow{p} \boldsymbol{\theta}^*$, we have $-\Psi_n'(\tilde{\boldsymbol{\theta}}) \xrightarrow{p} -\mathrm{E}\left[\psi_{\boldsymbol{\theta}}'(\boldsymbol{Z})\right]$.

Define $\boldsymbol{B} = \mathrm{E}\left[\psi_{\boldsymbol{\theta}}'(\boldsymbol{Z})\right]$ and $\boldsymbol{M} = \mathrm{Var}(\psi_{\boldsymbol{\theta}}(\boldsymbol{Z})) = \mathrm{E}\left[\psi_{\boldsymbol{\theta}}(\boldsymbol{Z})^{\otimes 2}\right]$. By Slutsky's theorem and the delta method, we have

$$\sqrt{n}\left(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\right) \Rightarrow \mathcal{N}\left(\mathbf{0}, \boldsymbol{B}^{-1}\boldsymbol{M}\boldsymbol{B}^{-1}\right). \tag{B.4}$$

This completes the proof.

B.2 Derivation of Sample Size Formulae

We derive the sample size formulae for comparing two DTRs which recommend different first-stage treatments that are embedded in a SMART in which a continuous repeated-measures outcome is collected throughout the study. These formulae are based on the regression analyses described in section 2.2 and a Wald test.

We consider a SMART in which the outcome is collected three timepoints: at baseline ($t = 0$), immediately before assessing response/non-response ($t = 1$), and at the end of the study ($t = 2$). We ignore the presence of baseline covariates $\boldsymbol{X}$ and assume $\boldsymbol{\mu}^{(d)}(\boldsymbol{\theta})$ is piecewise-linear in $\boldsymbol{\theta}$ (see, for example, model (2.1)).

Recall from section 2.3 that we wish to the null hypothesis $H_0 : \boldsymbol{c}^\top\boldsymbol{\theta} = 0$. In particular, we are interested in contrasts $\boldsymbol{c}$ which yield an end-of-study comparison between two embedded DTRs which recommend different first-stage treatments. Since a comparison of two embedded DTRs will yield a 1-degree of freedom Wald test, we use a Z statistic:

$$Z = \frac{\sqrt{n}\boldsymbol{c}^\top\hat{\boldsymbol{\theta}}}{\sigma_c},$$

where $\sigma_c = \sqrt{\boldsymbol{c}^\top\boldsymbol{B}^{-1}\boldsymbol{M}\boldsymbol{B}^{-1}\boldsymbol{c}}$. Under H_0, by asymptotic normality of $\sqrt{n}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})$, the test

statistic follows an asymptotic standard normal distribution. Suppose we wish to size the SMART to detect the alternative hypothesis $\boldsymbol{c}^\top \boldsymbol{\theta} = \Delta$. By the definition of type-II error, we have

$$\begin{aligned}
\beta &= P\left(\left|\frac{\sqrt{n}\boldsymbol{c}^\top \hat{\boldsymbol{\theta}}}{\sigma_c}\right| \leq z_{1-\alpha/2} \mid \boldsymbol{c}^\top \boldsymbol{\theta} = \Delta\right) \\
&= P\left(-z_{1-\alpha/2} \leq \frac{\sqrt{n}}{\sigma_c}\boldsymbol{c}^\top \hat{\boldsymbol{\theta}} \leq z_{1-\alpha/2} \mid \boldsymbol{c}^\top \boldsymbol{\theta} = \Delta\right) \\
&= P\left(-z_{1-\alpha/2} - \frac{\sqrt{n}}{\sigma_c}\boldsymbol{c}^\top \boldsymbol{\theta} \leq \frac{\sqrt{n}}{\sigma_c}\boldsymbol{c}^\top \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\right) \leq z_{1-\alpha/2} - \frac{\sqrt{n}}{\sigma_c}\boldsymbol{c}^\top \boldsymbol{\theta} \mid \boldsymbol{c}^\top \boldsymbol{\theta} = \Delta\right) \\
&= P\left(-z_{1-\alpha/2} - \frac{\sqrt{n}}{\sigma_c}\Delta \leq \frac{\sqrt{n}}{\sigma_c}\boldsymbol{c}^\top \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\right) \leq z_{1-\alpha/2} - \frac{\sqrt{n}}{\sigma_c}\Delta\right) \\
&= \Phi\left(z_{1-\alpha/2} - \frac{\sqrt{n}}{\sigma_c}\Delta\right) - \Phi\left(-z_{1-\alpha/2} - \frac{\sqrt{n}}{\sigma_c}\Delta\right) \\
&\leq \Phi\left(z_{1-\alpha/2} - \frac{\sqrt{n}}{\sigma_c}\Delta\right),
\end{aligned}$$

we arrive at the following form for the minimum-required sample size:

$$n \geq \left(z_{1-\alpha/2} + z_{1-\beta}\right)^2 \frac{\sigma_c^2}{\Delta^2}, \tag{B.5}$$

where z_p is the pth quantile of the standard normal distribution. Formula (B.5) is a fairly standard result in the clinical trials literature (Lachin 1981; Friedman, Furberg, and DeMets 2010); however, because of the dependence on σ_c, the formula is not useful as written. The goal of this appendix is to derive a closed-form upper bound on σ_c so as to obtain a sample size formula in terms of marginal quantities which can be more easily elicited from clinicians, or estimated from the literature.

Recall the definitions of $\boldsymbol{B}$ and $\boldsymbol{M}$ in equations (2.7) and (2.8), respectively. These quantities depend on $\boldsymbol{D}^{(d)}$, the partial derivative matrix of $\boldsymbol{\mu}^{(d)}(\boldsymbol{\theta})$ and $\boldsymbol{V}^{(d)}(\boldsymbol{\tau})$, the working covariance matrix for $\boldsymbol{Y}$. By assumed linearity of $\boldsymbol{\mu}^{(d)}(\boldsymbol{\theta})$, $\boldsymbol{D}^{(d)}$ is a fixed, constant matrix for all d. Furthermore, we assume that the working covariance matrix $\boldsymbol{V}^{(d)}(\boldsymbol{\tau})$ is correctly specified and satisifies working assumption A2 so that $\boldsymbol{V}^{(d)}(\boldsymbol{\tau}) = \boldsymbol{\Sigma}$ for all $d \in \mathscr{D}$. Note that $\boldsymbol{\Sigma}$ is non-random.

The estimand in equation (2.11) is a function of potential outcomes; as written in equations (2.7) and (2.8), $\boldsymbol{B}$ and $\boldsymbol{M}$ are functions of observed data. We begin by expressing $\boldsymbol{B}$

in terms of potential outcomes. Under the positivity, consistency, and sequential ignorability conditions (identifiability assumptions I1 to I3), we can apply lemma 4.1 of Murphy et al. (2001) so that

$$\begin{aligned}\boldsymbol{B} &= \sum_{d\in\mathscr{D}} \mathop{\mathrm{E}}_{A_1,R,A_2}\left[W^{(d)}(A_1,R,A_2)\boldsymbol{D}^{(d)}\left(\boldsymbol{V}^{(d)}(\boldsymbol{\tau})\right)^{-1}\left(\boldsymbol{D}^{(d)}\right)^{\top}\right] \\ &= \sum_{d\in\mathscr{D}} \boldsymbol{D}^{(d)}\boldsymbol{\Sigma}^{-1}\left(\boldsymbol{D}^{(d)}\right)^{\top}. \end{aligned} \tag{B.6}$$

We now turn our attention to $\boldsymbol{M}$. Expanding the outer product inside the expectation, we have

$$\begin{aligned}\boldsymbol{M} &= \mathop{\mathrm{E}}_{A_1,R,A_2,\boldsymbol{Y}}\left[\left(\sum_{d\in\mathscr{D}} W^{(d)}(A_1,R,A_2)\boldsymbol{D}^{(d)}\left(\boldsymbol{V}^{(d)}(\boldsymbol{\tau})\right)^{-1}\left(\boldsymbol{Y}-\boldsymbol{\mu}^{(d)}(\boldsymbol{\theta})\right)\right)^{\otimes 2}\right] \\ &= \sum_{d\in\mathscr{D}} \mathop{\mathrm{E}}_{A_1,R,A_2,\boldsymbol{Y}}\left[\left(W^{(d)}(A_1,R,A_2)\right)^2\left(\boldsymbol{D}^{(d)}\boldsymbol{\Sigma}^{-1}\left(\boldsymbol{Y}-\boldsymbol{\mu}^{(d)}(\boldsymbol{\theta})\right)\right)^{\otimes 2}\right] \\ &\quad + \sum_{d\neq d'} \mathop{\mathrm{E}}_{A_1,R,A_2,\boldsymbol{Y}}\Big[W^{(d)}(A_1,R,A_2)W^{(d')}(A_1,R,A_2)\boldsymbol{D}^{(d)}\boldsymbol{\Sigma}^{-1} \\ &\qquad\qquad \left(\boldsymbol{Y}-\boldsymbol{\mu}^{(d)}(\boldsymbol{\theta})\right)\left(\boldsymbol{Y}-\boldsymbol{\mu}^{(d')}(\boldsymbol{\theta})\right)^{\top}\boldsymbol{\Sigma}^{-1}\left(\boldsymbol{D}^{(d')}\right)^{\top}\Big]. \end{aligned} \tag{B.7}$$

Consider a single summand of the first term in equation (B.7). We can write this as

$$\boldsymbol{D}^{(d)}\boldsymbol{\Sigma}^{-1}\mathop{\mathrm{E}}_{A_1,R,A_2,\boldsymbol{Y}}\left[W^{(d)}(A_1,R,A_2)^2\left(\boldsymbol{Y}-\boldsymbol{\mu}^{(d)}(\boldsymbol{\theta})\right)^{\otimes 2}\right]\boldsymbol{\Sigma}^{-1}\left(\boldsymbol{D}^{(d)}\right)^{\top}$$

The inner expectation is a $T\times T$ matrix, the (i,j)th element of which is

$$\mathop{\mathrm{E}}_{A_1,R,A_2,\boldsymbol{Y}}\left[W^{(d)}(A_1,R,A_2)^2\left(Y_{t_i}-\mu^{(d)}_{t_i}(\boldsymbol{\theta})\right)\left(Y_{t_j}-\mu^{(d)}_{t_j}(\boldsymbol{\theta})\right)\right]. \tag{B.8}$$

Notice that the work above is design-independent: $\boldsymbol{B}$ and $\boldsymbol{M}$ have the same form as equations (B.6) and (B.7), respectively, for all designs. Below, we proceed only for design II, but derivations for designs I and III are analogous, substituting appropriate definitions of $W^{(d)}(A_1,R,A_2)$. Recall that, for design II, when all randomization probabilities are 0.5, $W^{(d)}(A_1,R,A_2) = 2\mathbb{1}_{\left\{A_1=a_1^{(d)}\right\}}(R+2(1-R)\mathbb{1}_{\left\{A_2=a_2^{(d)}\right\}})$. Further, we restrict our focus to three timepoints, denoted t_0 (baseline), $t_1=t^*$, and $t_2>t^*$.

Consider, for example, $t=t_1$. By repeated use of iterated expectation and application

of identifiability assumptions I2 and I3, equation (B.8) becomes

$$
\begin{aligned}
&\underset{Y_{t_0},A_1,Y_{t_1},R,A_2,Y_{t_2}}{\mathrm{E}}\left[W^{(d)}(A_1,R,A_2)^2\left(Y_{t_1}-\mu_{t_1}^{(d)}(\boldsymbol{\theta})\right)^2\right]\\
&=\underset{Y_{t_0},A_1,Y_{t_1},R,A_2}{\mathrm{E}}\left[4\mathbb{1}_{\left\{A_1=a_1^{(d)}\right\}}\left(R+4(1-R)\mathbb{1}_{\left\{A_2=a_2^{(d)}\right\}}\right)\left(Y_{t_1}-\mu_{t_1}^{(d)}(\boldsymbol{\theta})\right)^2\right]\\
&=\underset{Y_{t_0}^{(d)},A_1,Y_{t_1},R^{(a_1)},A_2^{(d)}}{\mathrm{E}}\left[4\mathbb{1}_{\left\{A_1=a_1^{(d)}\right\}}\left(R^{(a_1)}+4(1-R^{(a_1)})\mathbb{1}_{\left\{A_2=a_2^{(d)}\right\}}\right)\left(Y_{t_1}^{(d)}-\mu_{t_1}^{(d)}(\boldsymbol{\theta})\right)^2\right]\\
&=\underset{S_2(\bar{A}_1)}{\mathrm{E}}\left[4\mathbb{1}_{\left\{A_1=a_1^{(d)}\right\}}\left(R^{(a_1)}+4(1-R^{(a_1)})\underset{A_2|S_2(\bar{A}_1)}{\mathrm{E}}\left[\mathbb{1}_{\left\{A_2=a_2^{(d)}\right\}}\right]\right)\left(Y_{t_1}^{(d)}-\mu_{t_1}^{(d)}(\boldsymbol{\theta})\right)^2\right]\\
&=\underset{Y_{t_0}^{(d)},A_1,Y_{t_1}^{(d)},R^{(a_1)}}{\mathrm{E}}\left[4\mathbb{1}_{\left\{A_1=a_1^{(d)}\right\}}\left(2-R^{(a_1)}\right)\left(Y_{t_1}^{(d)}-\mu_{t_1}^{(d)}(\boldsymbol{\theta})\right)^2\right]. \qquad \text{(B.9)}\\
&=4\underset{Y_1^{(d)}}{\mathrm{E}}\left[\left(Y_1-\mu_1^{(d)}\right)^2\right]-2\underset{Y_1^{(d)},R^{(a_1)}}{\mathrm{E}}\left[\left(Y_1-\mu_1^{(d)}\right)^2R^{(a_1)}\right] \qquad \text{(B.10)}\\
&=4\sigma^2-2\,\mathrm{Cov}\left(\left(Y_1-\mu_1^{(d)}\right)^2,R^{(a_1)}\right)-2\,\mathrm{E}\left[R^{(a_1)}\right]\mathrm{E}\left[\left(Y_1-\mu_1^{(d)}\right)^2\right] \qquad \text{(B.11)}\\
&=2(2-r_{a_1})\sigma^2. \qquad \text{(B.12)}
\end{aligned}
$$

Equation (B.10) follows from equation (B.9) by identifiability assumption I3 and smoothing over $Y_{t_0}^{(d)}$, equation (B.11) arises from the definition of covariance, and equation (B.12) is a consequence of working assumption A1(b).

Similar derivations and applications of the remaining working assumptions allow us to bound $\boldsymbol{c}^\top\boldsymbol{B}^{-1}\boldsymbol{M}\boldsymbol{B}^{-1}\boldsymbol{c}$ above by

$$
\begin{aligned}
\boldsymbol{c}^\top\boldsymbol{B}^{-1}\boldsymbol{M}\boldsymbol{B}^{-1}\boldsymbol{c}&\leq 2\cdot\frac{1}{2}\left((2-r_1)+(2-r_{-1})\right)\boldsymbol{c}^\top\boldsymbol{B}^{-1}\left(\sum_{d\in\mathscr{D}}\boldsymbol{D}^{(d)}\boldsymbol{\Sigma}^{-1}\right)^{\otimes 2}\boldsymbol{B}^{-1}\boldsymbol{c}\\
&=\frac{4\sigma^2(1-\rho)\left(\rho^2+4\rho-\frac{1}{2}(r_1+r_{-1})(2\rho+1)+2\right)}{1+\rho}. \qquad \text{(B.13)}
\end{aligned}
$$

Plugging equation (B.13) into formula (B.5) leads to the aforementioned "sharp" sample size formula for design II. Some algebra shows that

$$
\sigma_c^2\leq 4\sigma^2\cdot\left(1-\rho^2\right)\cdot\frac{1}{2}\left((2-r_1)+(2-r_{-1})\right), \qquad \text{(B.14)}
$$

which allows for an easy-to-understand sample size formula. Plugging this result into formula (B.5), we arrive at formula (2.13).

BIBLIOGRAPHY

Almirall, D., C. DiStefano, Y.-C. Chang, S. Shire, A. Kaiser, X. Lu, I. Nahum-Shani, R. Landa, P. Mathy, and C. Kasari. 2016. "Longitudinal Effects of Adaptive Interventions With a Speech-Generating Device in Minimally Verbal Children With ASD." *J. Clin. Child Adolesc. Psychol.* 45 (4): 442–456. ISSN: 1537-4416, 1537-4424. doi:10.1080/15374416.2016.1138407.

Almirall, D., I. Nahum-Shani, N. E. Sherwood, and S. A. Murphy. 2014. "Introduction to SMART Designs for the Development of Adaptive Interventions: With Application to Weight Loss Research." *Transl. Behav. Med.* 4 (3): 260–274. ISSN: 1869-6716; EN :1613-9860. doi:10.1007/s13142-014-0265-0.

Almirall, D., I. Nahum-Shani, L. Wang, and C. Kasari. 2018. "Experimental Designs for Research on Adaptive Interventions: Singly and Sequentially Randomized Trials." In *Optimization of Behavioral, Biobehavioral, and Biomedical Interventions: Advanced Topics,* edited by L. M. Collins and K. C. Kugler, 89–120. Statistics for Social and Behavioral Sciences. Cham: Springer International Publishing. ISBN: 978-3-319-91776-4. doi:10.1007/978-3-319-91776-4_4.

Arndt, S., R. Jorge, C. Turvey, and R. G. Robinson. 2000. "Adding Subjects or Adding Measurements: Which Increases the Precision of Longitudinal Research?" *Journal of Psychiatric Research* 34 (6): 449–455. ISSN: 0022-3956. doi:10/bjgwtn.

Auyeung, S. F., Q. Long, E. B. Royster, S. Murthy, M. D. McNutt, D. Lawson, A. Miller, A. Manatunga, and D. L. Musselman. 2009. "Sequential Multiple-Assignment Randomized Trial Design of Neurobehavioral Treatment for Patients with Metastatic Malignant Melanoma Undergoing High-Dose Interferon-Alpha Therapy." *Clin. Trials J. Soc. Clin. Trials* 6 (5): 480–490. ISSN: 1740-7745, 1740-7753. doi:10.1177/1740774509344633.

Budney, A. J. 2014. *Behavioral Treatment of Adolescent Substance Use.* https://clinicaltrials.gov/ct2/show/

Chakraborty, B., and E. E. M. Moodie. 2013. *Statistical Methods for Dynamic Treatment Regimes.* Statistics for Biology and Health. New York, NY: Springer New York. ISBN: 978-1-4614-7427-2. doi:10.1007/978-1-4614-7428-9.

Chang, W., J. Cheng, J. J. Allaire, Y. Xie, and J. McPherson. 2018. *Shiny: Web Application Framework for R.*

Cohen, J. 1988. *Statistical Power Analysis for the Behavioral Sciences.* 2nd ed. Hillsdale, N.J: L. Erlbaum Associates. ISBN: 978-0-8058-0283-2.

Cole, S. R., and M. A. Hernn. 2008. "Constructing Inverse Probability Weights for Marginal Structural Models." *Am. J. Epidemiol.* 168 (6): 656–664. ISSN: 00029262. doi:10.1093/aje/kwn164.

Cook, N. R., and J. H. Ware. 1983. "Design and Analysis Methods for Longitudinal Research." *Annu. Rev. Public Health* 4 (1): 1–23. ISSN: 0163-7525, 1545-2093. doi:10/br5tbh.

Diegidio, P., S. Hermiz, J. Hibbard, M. Kosorok, and C. S. Hultman. 2017. "Hypertrophic Burn Scar Research: From Quantitative Assessment to Designing Clinical Sequential Multiple Assignment Randomized Trials." *Clin. Plast. Surg.* 44 (4): 917–924. ISSN: 00941298. doi:10.1016/j.cps.2017.05.024.

Eckshtain, D. 2013. *Using SMART Experimental Design to Personalize Treatment for Child Depression.* https://clinicaltrials.gov/ct2/show/NCT01880814.

Fitzmaurice, G. M., N. M. Laird, and J. H. Ware. 2011. *Applied Longitudinal Analysis.* 2nd ed. Wiley series in probability and statistics. OCLC: ocn708358043. Hoboken, N.J: Wiley. ISBN: 978-0-470-38027-7.

Fitzsimons, H., M. Tuten, K. O'Grady, M. S. Chisolm, and H. E. Jones. 2015. "A Smart Design: Response to Reinforcement-Based Treatment Intensity among Pregnant, Drug-Dependent Women." *Drug and Alcohol Dependence* 156:e69. ISSN: 0376-8716. doi:10/gfn9pt.

Friedman, L. M., C. Furberg, and D. L. DeMets. 2010. *Fundamentals of Clinical Trials.* 4th ed. New York: Springer. ISBN: 978-1-4419-1585-6 978-1-4419-1586-3.

Fu, S. S., A. J. Rothman, D. M. Vock, B. Lindgren, D. Almirall, A. Begnaud, A. Melzer, et al. 2017. "Program for Lung Cancer Screening and Tobacco Cessation: Study Protocol of a Sequential, Multiple Assignment, Randomized Trial." *Contemp. Clin. Trials* 60:86–95. ISSN: 15517144. doi:10.1016/j.cct.2017.07.002.

Gail, M., and R. Simon. 1985. "Testing for Qualitative Interactions between Treatment Effects and Patient Subsets." *Biometrics* 41 (2): 361–372. ISSN: 0006-341X (Print) 0006-341X (Linking). doi:10.2307/2530862.

Hamaker, E. L., and M. Wichers. 2017. "No Time Like the Present: Discovering the Hidden Dynamics in Intensive Longitudinal Data." *Curr. Dir. Psychol. Sci.* 26 (1): 10–15. ISSN: 0963-7214, 1467-8721. doi:10/f9r4kc.

Hibbard, J. C., J. S. Friedstat, S. M. Thomas, R. E. Edkins, C. S. Hultman, and M. R. Kosorok. 2018. "LIBERTI: A SMART Study in Plastic Surgery." *Clin. Trials* 15 (3): 286–293. ISSN: 1740-7745, 1740-7753. doi:10.1177/1740774518762435.

Kasari, C., A. Kaiser, K. Goods, J. Nietfeld, P. Mathy, R. Landa, S. Murphy, and D. Almirall. 2014. "Communication Interventions for Minimally Verbal Children with Autism: A Sequential Multiple Assignment Randomized Trial." *J. Am. Acad. Child Adolesc. Psychiatry* 53 (6): 635–646. ISSN: 15275418. doi:10.1016/j.jaac.2014.01.019.

Keener, R. W. 2010. *Theoretical Statistics: Topics for a Core Course.* Springer texts in statistics. OCLC: ocn688477588. New York: Springer. ISBN: 978-0-387-93838-7.

Kidwell, K. M. 2014. “SMART Designs in Cancer Research: Past, Present, and Future.” *Clin. Trials J. Soc. Clin. Trials* 11 (4): 445–456. ISSN: 1740-7745, 1740-7753. doi:10.1177/1740774514525691.

Kidwell, K. M., N. J. Seewald, Q. Tran, C. Kasari, and D. Almirall. 2018. “Design and Analysis Considerations for Comparing Dynamic Treatment Regimens with Binary Outcomes from Sequential Multiple Assignment Randomized Trials.” *J. Appl. Stat.* 45 (9): 1628–1651. ISSN: 0266-4763, 1360-0532. doi:10.1080/02664763.2017.1386773.

Kilbourne, A. M., K. M. Abraham, D. E. Goodrich, N. W. Bowersox, D. Almirall, Z. Lai, and K. M. Nord. 2013. “Cluster Randomized Adaptive Implementation Trial Comparing a Standard versus Enhanced Implementation Intervention to Improve Uptake of an Effective Re-Engagement Program for Patients with Serious Mental Illness.” *Implement. Sci.* 8 (1). ISSN: 1748-5908. doi:10.1186/1748-5908-8-136.

Kosorok, M. R., and E. E. M. Moodie, eds. 2015. *Adaptive Treatment Strategies in Practice: Planning Trials and Analyzing Data for Personalized Medicine.* Philadelphia, PA: Society for Industrial and Applied Mathematics. ISBN: 978-1-61197-417-1. doi:10.1137/1.9781611974188.

Lachin, J. M. 1981. “Introduction to Sample Size Determination and Power Analysis for Clinical Trials.” *Control. Clin. Trials* 2 (2): 93–113. ISSN: 01972456. doi:10.1016/0197-2456(81)90001-5.

Lavori, P. W., and R. Dawson. 2000. “A Design for Testing Clinical Strategies: Biased Adaptive within-Subject Randomization.” *J. R. Stat. Soc. Ser. A Stat. Soc.* 163 (1): 29–38. doi:10.1111/1467-985X.00154.

———. 2004. “Dynamic Treatment Regimes: Practical Design Considerations.” *Clin. Trials* 1 (1): 9–20. ISSN: 17407745, 17407753. doi:10.1191/1740774S04cn002oa.

———. 2014. “Introduction to Dynamic Treatment Strategies and Sequential Multiple Assignment Randomization.” *Clin. Trials J. Soc. Clin. Trials* 11 (4): 393–399. ISSN: 1740-7745, 1740-7753. doi:10.1177/1740774514527651.

Lei, H., I. Nahum-Shani, K. Lynch, D. Oslin, and S. Murphy. 2012. “A "SMART" Design for Building Individualized Treatment Sequences.” *Annu. Rev. Clin. Psychol.* 8 (1): 21–48. ISSN: 1548-5943, 1548-5951. doi:10.1146/annurev-clinpsy-032511-143152.

Liang, K.-Y., and S. L. Zeger. 1986. “Longitudinal Data Analysis Using Generalized Linear Models.” *Biometrika* 73 (1): 13–22. doi:10.1093/BIOMET/73.1.13.

Lipsitz, S., G. Fitzmaurice, D. Sinha, N. Hevelone, J. Hu, and L. L. Nguyen. 2017. “One-Step Generalized Estimating Equations With Large Cluster Sizes.” *J. Comput. Graph. Stat.* 26 (3): 734–737. ISSN: 1061-8600, 1537-2715. doi:10/gfn5fd.

Longford, N. T. 1999. "Selection Bias and Treatment Heterogeneity in Clinical Trials." *Stat. Med.* 18 (12): 1467–1474. ISSN: 0277-6715, 1097-0258. doi:10.1002/(SICI)1097-0258(19990630)18:12<1467::AID-SIM149>3.0.CO;2-H.

Lu, X., I. Nahum-Shani, C. Kasari, K. G. Lynch, D. W. Oslin, W. E. Pelham, G. Fabiano, and D. Almirall. 2016. "Comparing Dynamic Treatment Regimes Using Repeated-Measures Outcomes: Modeling Considerations in SMART Studies." *Stat. Med.* 35 (10): 1595–1615. doi:10.1002/sim.6819.

Markowitz, J. C., and B. L. Milrod. 2015. "What to Do When a Psychotherapy Fails." *Lancet Psychiatry* 2 (2): 186–190. ISSN: 22150366. doi:10/gfj39f.

Maxwell, S. E. 1998. "Longitudinal Designs in Randomized Group Comparisons: When Will Intermediate Observations Increase Statistical Power?" *Psychological Methods* 3 (3): 275–290. ISSN: 1082-989X. doi:10/dw3ktb.

McKay, J. R., M. L. Drapkin, D. H. A. Van Horn, K. G. Lynch, D. W. Oslin, D. DePhilippis, M. Ivey, and J. S. Cacciola. 2015. "Effect of Patient Choice in an Adaptive Sequential Randomization Trial of Treatment for Alcohol and Cocaine Dependence." *J. Consult. Clin. Psychol.* 83 (6): 1021–1032. ISSN: 1939-2117, 0022-006X. doi:10.1037/a0039534.

Murphy, S. A., M. J. van der Laan, J. M. Robins, Conduct Problems Prevention Research Group, and Group. 2001. "Marginal Mean Models for Dynamic Regimes." *J. Am. Stat. Assoc.* 96 (456): 1410–1423. ISSN: 0162-1459. doi:10.1198/016214501753382327.

Murphy, S. A. 2005. "An Experimental Design for the Development of Adaptive Treatment Strategies." *Stat. Med.* 24 (10): 1455–1481. ISSN: 0277-6715. doi:10.1002/sim.2022.

Murphy, S. A., K. G. Lynch, D. Oslin, J. R. McKay, and T. TenHave. 2007. "Developing Adaptive Treatment Strategies in Substance Abuse Research." *Drug Alcohol Depend.* 88:S24–S30. ISSN: 03768716. doi:10.1016/j.drugalcdep.2006.09.008.

Nahum-Shani, I., M. Qian, D. Almirall, W. E. Pelham, B. Gnagy, G. A. Fabiano, J. G. Waxmonsky, J. Yu, and S. A. Murphy. 2012a. "Experimental Design and Primary Data Analysis Methods for Comparing Adaptive Interventions." *Psychol. Methods* 17 (4): 457–477. ISSN: 1939-1463, 1082-989X. doi:10.1037/a0029372.

———. 2012b. "Q-Learning: A Data Analysis Method for Constructing Adaptive Interventions." *Psychol. Methods* 17 (4): 478–494. ISSN: 1082-989X. doi:10.1037/a0029373.

NeCamp, T., A. Kilbourne, and D. Almirall. 2017. "Comparing Cluster-Level Dynamic Treatment Regimens Using Sequential, Multiple Assignment, Randomized Trials: Regression Estimation and Sample Size Considerations." *Stat. Methods Med. Res.* 26 (4): 1572–1589. ISSN: 0962-2802, 1477-0334. doi:10.1177/0962280217708654.

Oetting, A. I., J. A. Levy, R. D. Weiss, and S. A. Murphy. 2011. "Statistical Methodology for a SMART Design in the Development of Adaptive Treatment Strategies." In *Causality and Psychopathology: Finding the Determinants of Disorders and Their Cures,* edited by P. E. Shrout, K. M. Keyes, and K. Ornstein, 179–205. New York: Oxford University Press. ISBN: 978-0-19-975464-9.

Ogbagaber, S. B., J. Karp, and A. S. Wahed. 2016. "Design of Sequentially Randomized Trials for Testing Adaptive Treatment Strategies." *Stat. Med.* 35 (6): 840–858. doi:10.1002/sim.6747.

Oslin, D. 2005. *Managing Alcoholism in People Who Do Not Respond to Naltrexone.* https://clinicaltrials.gov/ct2/show/NCT00115037.

Overall, J. E., and S. R. Doyle. 1994. "Estimating Sample Sizes for Repeated Measurement Designs." *Control. Clin. Trials* 15 (2): 100–123. ISSN: 01972456. doi:10.1016/0197-2456(94)90015-9.

Pelham, W. E., Jr., G. A. Fabiano, J. G. Waxmonsky, A. R. Greiner, E. M. Gnagy, W. E. P. III, S. Coxe, et al. 2016. "Treatment Sequencing for Childhood ADHD: A Multiple-Randomization Study of Adaptive Medication and Behavioral Interventions." *J. Clin. Child Adolesc. Psychol.* 45 (4): 396–415. ISSN: 1537-4416. doi:10/gfn9xr.

Raudenbush, S. W., and L. Xiao-Feng. 2001. "Effects of Study Duration, Frequency of Observation, and Sample Size on Power in Studies of Group Differences in Polynomial Change." *Psychol. Methods* 6 (4): 387–401. ISSN: 1082-989X. doi:10.1037//1082-989X.6.4.387.

Robins, J. M. 1997. "Causal Inference from Complex Longitudinal Data." In *Latent Variable Modeling and Applications to Causality,* edited by P. Bickel, P. Diggle, S. Fienberg, K. Krickeberg, I. Olkin, N. Wermuth, S. Zeger, and M. Berkane, 120:69–117. New York, NY: Springer New York. ISBN: 978-0-387-94917-8 978-1-4612-1842-5. doi:10.1007/978-1-4612-1842-5_4.

Thall, P. F. 2015. "SMART Design, Conduct, and Analysis in Oncology." In *Adaptive Treatment Strategies in Practice: Planning Trials and Analyzing Data for Personalized Medicine,* edited by M. R. Kosorok and E. E. M. Moodie, 41–54. Philadelphia, PA: Society for Industrial and Applied Mathematics. ISBN: 978-1-61197-417-1.

van der Vaart, A. W. 1998. *Asymptotic Statistics.* Cambridge series in statistical and probabilistic mathematics. Cambridge, UK ; New York, NY, USA: Cambridge University Press. ISBN: 978-0-521-49603-2.

Wallace, M. P., and E. E. M. Moodie. 2014. "Personalizing Medicine: A Review of Adaptive Treatment Strategies." *Pharmacoepidemiol. Drug Saf.* 23 (6): 580–585. ISSN: 1099-1557 (Electronic)\r1053-8569 (Linking). doi:10.1002/pds.3606.

Walls, T. A., and J. L. Schafer, eds. 2006. *Models for Intensive Longitudinal Data.* OCLC: ocm58720696. Oxford ; New York: Oxford University Press. ISBN: 978-0-19-517344-4.

Watkins, C. J. C. H. 1989. "Learning from Delayed Rewards." PhD diss., King's College.

Zhang, S., and C. Ahn. 2011a. "Adding Subjects or Adding Measurements in Repeated Measurement Studies Under Financial Constraints." *Stat. Biopharm. Res.* 3 (1): 54–64. ISSN: null. doi:10/bftphq.

———. 2011b. "How Many Measurements for Time-Averaged Differences in Repeated Measurement Studies?" *Contemporary Clinical Trials* 32 (3): 412–417. ISSN: 1551-7144. doi:10/dwv3qn.

Zhang, S., J. Cao, and C. Ahn. 2014. "A GEE Approach to Determine Sample Size for Pre- and Post-Intervention Experiments with Dropout." *Comput. Stat. Data Anal.* 69:114–121. ISSN: 01679473. doi:10.1016/j.csda.2013.07.037.

Zhang, Y., E. B. Laber, A. Tsiatis, and M. Davidian. 2015. "Using Decision Lists to Construct Interpretable and Parsimonious Treatment Regimes." *Biometrics* 71 (4): 895–904. ISSN: 0006341X. doi:10.1111/biom.12354.