DOUBLY ROBUST ESTIMATION IN TWO-STAGE DYNAMIC TREATMENT REGIMES IN THE PRESENCE OF DROP-OUT

by

Andrew S. Topp

B.S. Mathematics, University of Arizona, 2011

Submitted to the Graduate Faculty of

the Graduate School Of Public Health in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2016

UNIVERSITY OF PITTSBURGH

GRADUATE SCHOOL OF PUBLIC HEALTH

This dissertation was presented

by

Andrew S. Topp

It was defended on

September 20th 2016

and approved by

Abdus S. Wahed, PhD

Professor

Department of Biostatistics

Graduate School Of Public Health

University of Pittsburgh

Committee Member:

Sally Morton, PhD

Department Chair

Department of Biostatistics

Graduate School Of Public Health

University of Pittsburgh

Committee Member:

Doug Landsittel, PhD

Professor

Department of Medicine

School of Medicine

University of Pittsburgh

Committee Member:

Stewart Anderson, PhD

Professor

Department of Biostatistics

Graduate School Of Public Health

University of Pittsburgh

Committee Member:

Chaeryon Kang, PhD

Assistant Professor

Department of Biostatistics

Graduate School Of Public Health

University of Pittsburgh

Dissertation Director: Abdus S. Wahed, PhD

Professor

Department of Biostatistics

Graduate School Of Public Health

University of Pittsburgh

Abdus S. Wahed, PhD

DOUBLY ROBUST ESTIMATION IN TWO-STAGE DYNAMIC TREATMENT REGIMES IN THE PRESENCE OF DROP-OUT

Andrew S. Topp, PhD

University of Pittsburgh, 2016

Abstract

Various methods exist for causal inference about the effects of different treatments from observational studies or randomized trials. A straightforward approach is to fit a regression model of the outcome as a function of the treatment they received along with observed patient characteristics. Other methods, such as inverse probability weighting, work by instead estimating a patient's probability of receiving treatment and weighting the outcomes by the inverse of this probability. Doubly Robust estimators use both models and provide unbiased estimates as long as either the probability of treatment or outcome is correctly modeled. These techniques can be extended to analyze and compare dynamic treatment regimes, that is, multiple stages of treatment punctuated by decision points concerning what the next treatment should be. Observational studies and randomized trials represent a cornerstone of public health research. As they grow and become more complex it is important for the wealth of statistical methods to grow with them.

This dissertation is concerned with developing more efficient doubly robust estimators for two-stage dynamic treatment regimes, first without and later, with data missing at random. First, we develop a new inverse probability of treatment weighted and doubly robust estimators for analyzing dynamic treatment regimes. Then we compare these methods to the corresponding existing methods for estimating the mean outcome of dynamic treatment regimes in a simulation. We utilize the new doubly robust estimator in the analysis of the STAR*D trial to estimate the mean outcome of patients on different regimes for the treatment of non-psychotic major depressive disorder. Finally we propose a modification of the new inverse probability of treatment weighted and doubly robust estimators in order to account for missing data.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

PREFACE

I would first and foremost like to express my gratitude to my advisor Dr. Abdus S Wahed. He has been my guide through the majority of my graduate school experience and has shown more patience working with me than I thought any one person was capable of possessing. Thank you for allowing me to become your dissertation advisee, your abilities as a teacher and mentor are unparalleled.

The faculty of the Department of Biostatistics at the University of Pittsburgh has continued to surprise me over the course of the past five years in how supportive and nuturing they managed to be, even when I least expected it. Thank you to Dr. Sally Morton for not only being an excellent department head but an academic advisor who guided me through my first few years of school with great advice, always managing to point me in the right direction. Thank you to Dr. Doug Landsittel for supporting me as a graduate student researcher (GSR). Instead of treating me as only an employee you expressed an interest and invested time in my growth as a statistician, finding time to talk to me about subjects outside of my duties as a GSR despite your overwhelming schedule. Thank you also to Dr. Geoff Johnson, you deserve thanks for so much from the time we spent together at Pitt. However, I've decided to use this space to call you out for throwing paper airplanes during Dr. Anderson's (excellent) class.

I dedicate this work to everyone who helped carry me this far. I owe a great debt of gratitude to all my friends, loved ones, mentors, and teachers that have inspired, supported, and stood by me these past five years.

What follows is the not just the fruit of my labor but the result of the guidance and input of a great many people. Each chapter is designed to stand on its own so some information is repeated between chapters.

1.0 INTRODUCTION

In statistics, we are often interested in finding and estimating causal relationships between variables. Challenges such as confounding, unobserved data, and lack of prior knowledge hinder our ability to make accurate and useful causal inference about the relationships between these variables. Recent history has seen increased interest in establishing rigorous methodology for the purpose of making causal inference.

1.1 ASSUMPTIONS OF CAUSAL INFERENCE

The first assumption required for causal inference is the consistency assumption (Robins et. al [2000]), which means that the outcome a patient experiences under a treatment is in fact, the outcome caused by receiving that treatment. For patients that do not receive the treatment of interest, that quantity is not observed. Additionally, causal inference requires that every combination of treatments be possible for all subjects, this is known as the positivity assumption [Hernan and Robins 2006]. Lastly, it needs the sequential randomization assumption [Robins 1986], which states that given a subject's covariate history and prior treatment assignment, the treatment assignment at any given decision point is made independent of all potential

outcomes. This assumption is known as the no unmeasured confounder assumption. These assumptions are necessary for unbiased, or at least, consistent, estimation of treatment effects and the corresponding statistical inference.

1.2 MODELING OUTCOME AND G-COMPUTATION

Consider a trial looking at treatments A and B and comparing their effects on patient's outcome Y. For each patient we can only see their outcome given treatment for one of the two treatments, i.e., for a patient who receives treatment A, we see that patient's outcome of Y given A. However we do not observe Y given B, i.e. patient's outcome had he or she received treatment B. We refer to this unobserved potential outcome as a counterfactual outcome.

While counterfactual variables are not always observable they provide avenues to construct valid estimates and to assess treatment effects from randomized trials or observational studies. Various methods have been proposed that produce unbiased and consistent estimates of treatment effects based on the use of counterfactuals. The standard and most straightforward method is to directly model the outcome given receiving each treatment. However, the relationship between the observed covariates and the outcome may be unknown. In the case of non-randomized studies, there may be unobserved confounders important to the causal relationship between treatment and outcome.

When comparing treatments within a study, investigators can only observe the outcome corresponding to the treatment the patient was assigned. This is analogous to the missing data problem frequently encountered in statistical analysis of biomedical data. All the unobserved potential outcomes can be thought of as missing. One

common approach to dealing with missing data is to replace the missing values with approximations based on the observed data. In regression-based imputation, we construct a model of the outcome regressed on treatment and any relevant covariates, which allows us to estimate the conditional expectation of the outcome given the treatment the patient did not receive using the resulting model. An unbiased estimate of the mean outcome conditional on receiving a particular treatment is then constructed by these conditional expectations. More specifically, this estimate of the mean outcome can be expressed as $\hat{\mu}_{OR} = \frac{1}{n}\sum_{i=1}^{n} m(X_i, \hat{\beta})$, where $m(X_i, \beta)$ is the postulated outcome model for $E[Y_i|X_i]$, n is the number of subjects in the sample, X_i are the covariates of interest, and $\hat{\beta}$ is a consistent estimator of β, the parameter for the outcome model. However, if we misspecify the outcome model $m(X_i, \beta)$ we no longer get a unbiased and consistent estimator of the mean outcome under our treatment of interest.

G-computation is an extension of ordinary least squares regression in settings where the population is stratified by intermediate responses to treatments. The G-computation algorithm was developed by Robins [1986] and [1987] in order to conduct causal inference when patients experience a sustained exposure period. It has since been used in the analysis of multi-stage studies (see for example, Bembom and Van der Laan [2007]). In multi-stage treatment settings, the algorithm works by estimating the distribution of the intermediate outcomes between stages of treatment and uses weighted averages to obtain estimates of population means.

1.3 INVERSE PROBABILITY OF TREATMENT WEIGHTING

One approach to causal inference analysis is called inverse probability of treatment weighting (IPTW). In this approach each observation is weighted with the inverse of the probability of receiving the treatment the patient received. So a patient who received treatment A would have a weight of $1/pr(Receiving\ A)$. With this method, a patient counts for herself along with $1/pr(Receiving\ A) - 1$ other patients that have similar covariate values but for whom we did not see the counterfactual outcome given treatment A, because they received other treatments. In classical randomized trials, these probabilities are known and hence it would be straightforward to apply inverse probability of treatment weighting. In observational studies, however, treatment assignment probabilities are not known, and need to be estimated from the data. If we model the probability of receiving treatment based on the observed covariates the analysis may be flawed in the same way as when modeling the outcome incorrectly in the previous method. In this case, the accuracy of the estimate will depend on the relationship between the probability of receiving treatment and the included covariates. Incorrect specification of this relationship could lead to biased and/or inefficient inference. Using this method to estimate the mean counterfactual outcome of the study population receiving A, which we define as $Y[A]$, would result in the Inverse Probability of Treatment Weighting (IPTW) estimator [Horvitz and Thompson, 1952]

$$\hat{\mu}_{IPTW} = \frac{1}{n}\sum_{i=1}^{n} \frac{Z_i^{(A)}}{\pi_i^{(A)}} Y_i \ , \tag{1.1}$$

where $Z_i^{(A)}$ is the indicator for receiving treatment A, $Z_i^{(A)} = 1$, if the patient receives A, 0, otherwise, $\pi_i^{(A)}$ is the probability of receiving treatment A, and Y_i is the outcome for the i^{th} subject, $i = 1, ..., n$. We can see that this estimator is unbiased under correct specification of the treatment model $\pi_i^{(A)}$ by taking its expectation.

$$\begin{aligned}
E[\hat{\mu}_{IPTW}] &= E\left[\frac{1}{n}\sum_{i=1}^{n}\frac{Z_i^{(A)}}{\pi_i^{(A)}}Y_i\right] = \frac{1}{n}\sum_{i=1}^{n}E\left[\frac{Z_i^{(A)}}{\pi_i^{(A)}}Y_i\right] \\
&= \frac{1}{n}\sum_{i=1}^{n}E\left[\frac{Z_i^{(A)}}{\pi_i^{(A)}}(Y_i[A]Z_i^{(A)} + Y_i[\neg A](1 - Z_i^{(A)}))\right] \\
&= \frac{1}{n}\sum_{i=1}^{n}E\left[\frac{Z_i^{(A)}}{\pi_i^{(A)}}Y_i[A]\right] = \frac{1}{n}\sum_{i=1}^{n}E\left[E\left[\frac{Z_i^{(A)}}{\pi_i^{(A)}}Y_i[A]\middle|Y_i[A]\right]\right] \\
&= \frac{1}{n}\sum_{i=1}^{n}E\left[Y_i[A]E\left[\frac{Z_i^{(A)}}{\pi_i^{(A)}}\middle|Y_i[A]\right]\right] = \frac{1}{n}\sum_{i=1}^{n}E[Y_i[A]] = E[Y[A]] \ ,
\end{aligned}$$

since $Pr(Z_i^{(A)} = 1) = \pi_i^{(A)}$ under correct treatment model specification.

The past few decades have seen a great expansion of the literature on inverse probability weighting. Books by van der Laan and Robins [2003] and Tsiatis [2006] discuss inverse probability weighted augmented estimation, that is, inverse probability weighted estimators with an additional augmentation term that models the coarsening mechanism . Robins and Rotnitzky [1992], Robins, Rotnitzky, and Zhao [1994], and Robins and Rotnitzky [1995] all discuss locally efficient inverse probability weighted augmented estimation. These estimators are referred to as locally efficient estimators because they satisfy the conditions that they are both asymptotically normal even when their working models are misspecified and that when they are correctly specified they have the smallest variance among all inverse probability weighted augmented estimators. Furthermore, Rotnitsky, Robins, and Scharfstein [1998], Robins, Rotnitsky, and Scharfstein [2000], Scharfstein, Rotnitzky, and Robins

[1999], Scharfstein and Irrizarry [2003], and Birmingham, Rotnitzky, and Fitzmaurice [2003] all worked on inverse probability weighted augmented estimation in situations where data are not missing completely at random. Lastly, Robins [1999] discusses using inverse probability weighting to derive parameter estimates in causal inference models.

1.4 DOUBLY ROBUST ESTIMATION

There are methods requiring the specification of both outcome and treatment models, called doubly-robust methods, that allow for unbiased or consistent estimates. This continues to be true if at least one of the models is correctly specified. Doubly robust estimation will be the focus of this paper. In the setting described in previous sections, the doubly robust estimator for mean outcome is

$$\hat{\mu}_{DR} = \frac{1}{n}\sum_{i=1}^{n} \left\{ \frac{Z_i^{(A)}}{\pi^{(A)}(X_i,\hat{\gamma})}\, Y_i - \frac{Z_i^{(A)} - \pi^{(A)}(X_i,\hat{\gamma})}{\pi^{(A)}(X_i,\hat{\gamma})} m(X_i,\hat{\beta}) \right\}, \qquad (1.2)$$

where $\pi^{(A)}(X_i,\hat{\gamma})$ and $m(X_i,\hat{\beta})$ are the i^{th} subject's estimated probability of treatment and outcome respectively. When the treatment model is true and the coefficient γ is known then $E\left[\frac{Z_i^{(A)}-\pi^{(A)}(X_i,\gamma)}{\pi^{(A)}(X_i,\gamma)}\middle| X_i\right] = \frac{P(Z_i^{(A)}=1|X_i)-\pi^{(A)}(X_i,\gamma)}{\pi^{(A)}(X_i,\gamma)} = 0$ and $E[\hat{\mu}_{DR}]$ becomes the expected value of the inverse probability of treatment weighted estimator, and hence is an unbiased estimator of μ. If γ is estimated using a correctly specified logistic regression model then the estimator is still consistent. Likewise, when the

outcome model is true and the coefficient β is known we can see

$$E[\hat{\mu}_{DR}] = E\left[\frac{1}{n}\sum_{i=1}^{n}\left\{\frac{Z_i^{(A)}}{\pi^{(A)}(X_i,\hat{\gamma})}\,Y_i - \frac{Z_i^{(A)}-\pi^{(A)}(X_i,\hat{\gamma})}{\pi^{(A)}(X_i,\hat{\gamma})}m(X_i,\beta)\right\}\right]$$
$$= \frac{1}{n}\sum_{i=1}^{n}\left\{\frac{E[Y_i[A]]E[Z_i^{(A)}|Y_i[A]]}{\pi^{(A)}(X_i,\hat{\gamma})} - \frac{E[Z_i^{(A)}|Y_i[A]]-\pi^{(A)}(X_i,\hat{\gamma})}{\pi^{(A)}(X_i,\hat{\gamma})}E[Y_i[A]]\right\} \quad (1.3)$$
$$= \frac{1}{n}\sum_{i=1}^{n}\left\{\frac{\pi^{(A)}(X_i,\hat{\gamma})}{\pi^{(A)}(X_i,\hat{\gamma})}E[Y_i[A]]\right\} = \frac{1}{n}\sum_{i=1}^{n}E[Y_i[A]] = E[Y[A]]\ ,$$

which is consistent when the βs are correctly estimated.

There has been a wealth of literature published on the subject of doubly robust estimation since its introduction by Robins [Robins, Am. Statist. Assoc. Sect. Bayesian Statist. Sci 1999]. Carpenter and Kenward [2006] compare using this method to the more common approach of multiple imputation. Both Bang and Robins [2005] and Tsiatis, Davidian, and Cao [2011] utilize doubly robust estimators in a longitudinal setting and improve its efficiency. Kang and Schafer [2007] compare various doubly robust, as well as some non-doubly robust, methods in different missing data scenarios. Rotnitzky, Lei, Sued, and Robins [2012] discuss a doubly robust estimator derived by solving outcome regression estimating equations. Bai, Tsiatis, and O'Brien [2013] examine a doubly robust estimator in the survival data setting, with particular emphasis on stratified sampling schemes.

Cao, Tsiatis, and Davidian [2009] in particular are concerned with the efficient estimation of the outcome model coefficients in cross-sectional settings. They observe that even if one correctly specified the treatment model with incorrect specification of the outcome model the results, while still yielding unbiased estimates, will be inefficient. They propose an estimating equation for these coefficients that results in more efficient doubly robust estimates.

1.5 DYNAMIC TREATMENT REGIMES

Sometimes chronic diseases demand frequent modification of treatments based on individual responses to the treatment and/or health conditions. Therefore, instead of caring about which treatment is immediately better a physician may want to know what course of treatments is best overall. These courses of treatments are called dynamic treatment regimes. They are rules of treatment choice based on intermediate responses and patient characteristics. For example, a dynamic treatment regime in depression treatment could be 'treat a patient suffering from clinical depression with a selective serotonin re-uptake inhibitor (SSRI), if they respond, keep them on it, and if not switch them to bupropion (a non-SSRI treatment for clinical depression).' G-computation, Inverse Probability of Treatment Weighting, and Doubly Robust estimation all have applications in analyzing and comparing dynamic treatment regimes.

Dynamic Treatment Regimes can be examined in a clinical trial setting using sequential, multiple assignment, randomized trial, or SMART, designs (Murphy et al. 2005). The idea behind these designs is to have multiple stages of treatment punctuated by randomization at decision points to treatments within certain treatment groups depending on their intermediate responses. For example, continuing with our clinical depression example, a SMART design would start patients off on an SSRI like citalopram and at the end of the stage decide if the patient had responded. Responders could be randomized to a maintenance treatment or simply kept on the medication that has proven to be effective while non-responders could be randomized to a different treatment altogether. A visualization of this design can be seen in Figure 1.

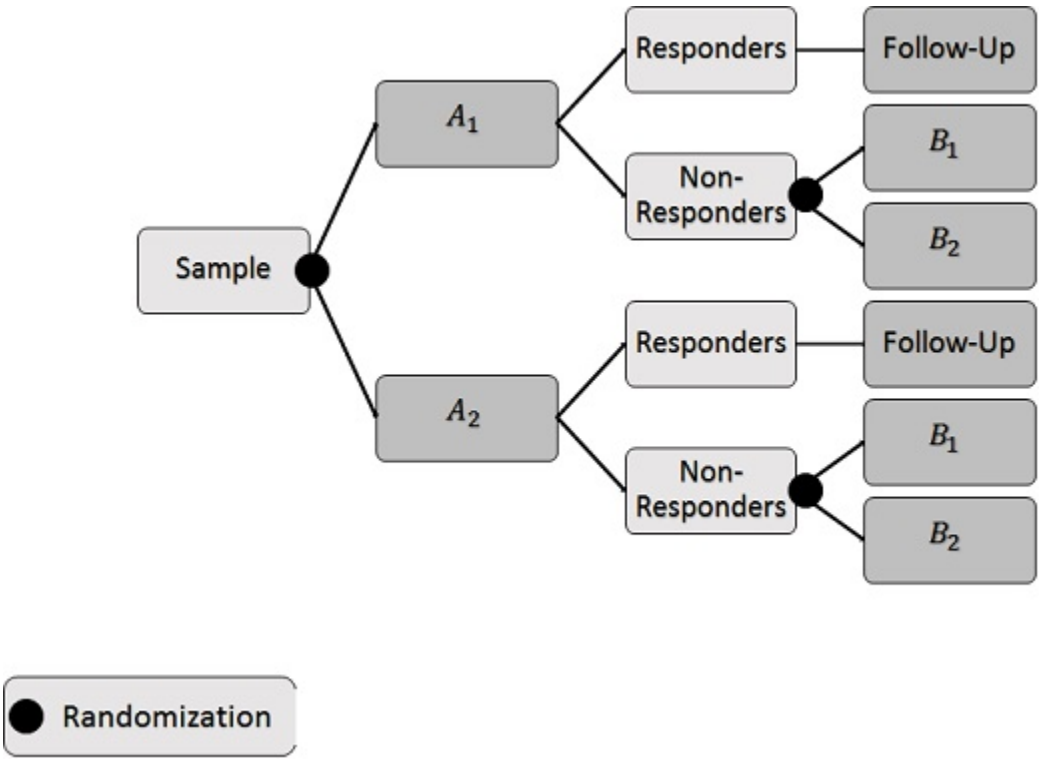

Figure 1: SMART Design; A_j = 1st stage Treatment, j=1,2, B_k = 2nd stage Treatment, k=1,2,

1.6 THE STAR*D TRIAL

The Sequenced Treatment Alternatives to Relieve Depression, or STAR*D, Trial was a multistage trial designed to test different regimes of depression treatments. The trial contained 4,041 patients between the ages of 18 and 75 who had received a clinical diagnosis of nonpsychotic major depressive disorder under the DSM-IV checklist. A patient that is considered to respond to treatment will continue taking that treatment, while those that do not respond are re-randomized to a treatment within a treatment-grouping of their own choice.

The trial was designed with four (and a half) stages. This study contains a unique element called Equipoise Stratified Randomization, meaning that patients were given the option of selecting classes of treatments that they would find acceptable to be on and then were randomized to one of the treatments within those groups. This was done in order to boost compliance and decrease drop-out. A patient with more control over their treatment is thought to be more inclined to follow that treatment and not leave the study. Everyone received the same treatment in the first stage, an SSRI called citalopram. Non-responders at each stage were offered a choice of switching off of their current medication or supplementing it with another treatment. Non-responders to the second stage who received cognitive therapy or citalopram supplemented with cognitive therapy that stage were randomized to one of two medicinal treatments at what investigators called stage 2A. The purpose of this was to make sure everyone that made it to the third stage received a follow up medication to citalopram.

Response between stages was defined as either a 50% reduction in QIDS-16 (Quick Inventory of Depressive Symptomatology) self reported score, which the investigators call response, or a QIDS-16 score lower than 5, which the investigators call remission,

at the end of the stage. The overall outcome of interest for investigators was the HRSD-17 (Hamilton Rating Scale of Depression) score which was collected upon study exit.

1.7 MOTIVATION

As mentioned there are a number of existing methods to estimate the mean outcomes of dynamic treatment regimes. The first method we described was G-computation which directly models outcome and intermediate response. Alternatively, Inverse Probability of Treatment Weighted (IPTW) estimators utilize the modeled probability of treatment as weights to account for the missing counterfactual outcomes of those that did not receive the regime of interest. Finally, Doubly Robust estimators allow for unbiased estimation of the mean regime outcomes as long as either the treatment or outcome models are correctly specified.

The existing doubly robust estimator for two-stage dynamic treatment regimes (Wahed and Tsiatis [2004]) utilizes inverse probability of treatment weighting to account for patients deviating from second stage treatments, but not first stage treatments, specified in the regime of interest. This means that the estimator does not take advantage of all the available data for patients with missing data. This estimator could be modified to use all the first stage treatment data. However, the estimator becomes inefficient and biased because the model based correction piece no longer has expectation zero under correct treatment model specification. This is no longer the case when the first stage treatment assignment and intermediate response are independent but independence is an unreasonable assumption to make.

Our goal is to build a new doubly robust estimator that can utilize inverse prob-

ability of treatment weighting for both the first and second stages of treatment while not requiring the same strong independence assumption as the modified existing estimator. We expect this estimator to be more efficient and less biased under model misspecification.

2.0 EFFICIENT DOUBLY ROBUST ESTIMATION FOR TWO-STAGE DYNAMIC TREATMENT REGIMES

2.1 INTRODUCTION

When comparing two or more treatments in observational or randomized studies, for a given patient, investigators can only observe the outcome corresponding to the treatment received by the patient. Thus, investigators are unable to observe the potential outcomes that could have been observed had a patient received other treatments. These potential outcomes a patient could have experienced are often referred to as counterfactual outcomes [Holland 1986].

While counterfactual variables are not always observable, they provide avenues to construct valid estimates and to assess treatment effects from randomized trials or observational studies. Various methods have been proposed that produce unbiased and consistent estimates of treatment effects based on the use of counterfactuals. The standard and most straightforward method is to directly model the outcome given receiving each treatment. However, modeling an outcome may not be straightforward because of the uncertainty of the relationship between the observed covariates and the outcome. In the case of non-randomized studies, there may be unobserved confounders important to the causal relationship between treatment and outcome.

An alternate approach is called inverse probability of treatment weighting. In this approach each observation is weighted with the inverse of the probability of receiving the treatment they received. So a patient who received treatment A_1 would have a weight of $1/pr(ReceivingA_1)$. With this method, a patient counts for herself along with $1/pr(ReceivingA_1) - 1$ other patients that have similar covariate values but for whom we did not see the counterfactual outcome given treatment A_1. In classical randomized trials, these probabilities are known and hence it would be straightforward to apply inverse probability of treatment weighting. In observational studies, however, treatment assignment probabilities are not known, and needs to be estimated from the data. If we model the probability of receiving treatment based on the observed covariates the analysis may be flawed in the same way as when modeling the outcome incorrectly in the previous method. In this case, the accuracy of the estimate will depend on the relationship between the probability of receiving treatment and the included covariates. Incorrect specification of this relationship could lead to biased and/or inefficient inference. However, there are methods requiring the specification of both outcome and treatment models, called doubly-robust methods, that allow for unbiased or consistent estimates. This continues to be true if at least one of the models is correctly specified. Doubly robust estimation will be the focus of this paper.

There has been a wealth of literature published on the subject of doubly robust estimation since its introduction by Robins [Robins Am. Statist. Assoc. Sect. Bayesian Statist. Sci 1999] . Carpenter and Kenward [2006] compare using this method to the more common approach of multiple imputation. Both Bang and Robins [2005] and Tsiatis, Davidian, and Cao [2011] utilize doubly robust estimators in a longitudinal setting and attempt to improve to this method's efficiency. Kang and Schafer [2007] compare various doubly robust, as well as some non-doubly robust

methods in different missing data scenarios. Rotnitzky, Lei, Sued, and Robins [2012] discuss a doubly robust estimator derived by solving outcome regression estimating equations. Bai, Tsiatis, and O'Brien [2013] examine a doubly robust estimator in the survival data setting, with particular emphasis on stratified sampling schemes.

Cao, Tsiatis, and Davidian [2009] in particular are concerned with the efficient estimation of the outcome model coefficients in cross-sectional settings. They observe that even if one correctly specified the treatment model with incorrect specification of the outcome model, the results, while still yielding unbiased estimates, will be inefficient. They propose an estimating equation for these coefficients that results in more efficient doubly robust estimates.

Doubly robust estimators discussed above were mostly derived or implemented for settings where one treatment is compared to another in cross-sectional or longitudinal settings, but chronic diseases demand frequent modification of treatments based on individual responses to the treatment and/or health conditions. Therefore, instead of caring about which treatment is immediately better, a physician may want to know what course of treatments is best overall. These courses of treatments are called dynamic treatment regimes (DTRs). They are rules of treatment choice based on intermediate responses and patient characteristics. For example, a dynamic treatment regime in depression treatment could be 'treat a patient suffering from clinical depression with a selective serotonin re-uptake inhibitor (SSRI), if they respond keep them on it, and if not switch them to a drug like bupropion (a non-SSRI treatment for clinical depression)'.

Comparing dynamic treatment regimes using counterfactual outcomes is generally more complicated than comparing a single set of treatments. For the latter it is enough to consider counterfactuals related to a group of treatments, but we need to think about all possible combinations of initial and follow-up treatments a

patient could experience in a dynamic treatment regime setting. Dynamic treat-ment regimes and related analytical techniques have grown in prominence in the past two decades. Murphy [2003] discusses the use of backwards induction and Q-learning as a technique for estimating regime outcomes. Bembom and van der Laan [2007] explore both inverse probability weighting and g-computation in two stage dynamic treatment regime clinical trials. Lunceford, Davidian, and Tsiatis [2002] discuss an inverse probability of treatment weighted estimator with a correctional constant chosen to minimize variance across a class of estimators they introduce. Zhang, Tsiatis, Laber, and Davidian [2013] propose a doubly robust estimator in a dynamic treatment regime setting whose efficiency improves through careful choice of an augmentation term.

Many well known, large studies aim to compare dynamic treatment regimes. The STAR*D trial [Warden et al. 2007], which motivated this research, is a multistage clinical trial examining the effects of depression treatments for patients who do not respond to SSRIs. The COG study A3891 [Matthay et al. 1999] examines the efficacy of chemotherapy followed by bone marrow transplantation versus a second round of chemotherapy in children with high-risk neuroblastoma. The REVAMP study [Trivedi et al. 2008] is another multistage study concerned with analyzing regimes of treatments for chronic depression.

The goal of this paper is to examine doubly robust estimators in a dynamic treatment regime setting in order to demonstrate a more efficient construction of the two-stage doubly robust estimator.

2.2 FRAMEWORK

2.2.1 Estimating Causal Treatment Effect

As mentioned in the section above when comparing treatments within a study, investigators can only observe the outcome corresponding to the treatment the patient was assigned. This is analogous to the missing data problem frequently encountered in statistical analysis of biomedical data. All the unobserved potential outcomes can be thought of as missing. One common approach to dealing with missing data is to replace the missing values with an approximation based on the observed data. In regression-based imputation, we construct a model of the outcome regressed on treatment and any relevant covariates, which allows us to estimate the conditional expectation of the outcome given the treatment the patient did not receive using the resulting model. An unbiased estimate of the mean outcome conditional on receiving a particular treatment is then constructed by these conditional expectations. More specifically $\hat{\mu}_{OR} = \frac{1}{n}\sum_{i=1}^{n} m(X_i, \hat{\beta})$, where $m(X_i, \beta)$ is the postulated ordinary least squares regression model for $E[Y_i|X_i]$, n is the number of subjects in the sample, and $\hat{\beta}$ is a consistent estimator of β, the parameter for the outcome model.

An alternative approach would be to impute the missing outcomes using the outcomes of other, similar patients. Assume a patient received treatment A with the probability of receiving it $\pi^{(A)}(X_i, \gamma)$, where γ are coefficients that describe the relationship between the covariates, X_i, and probability of receiving treatment. Then we can assume that there are $1/\pi^{(A)}(X_i, \gamma) - 1$ patients with covariates X_i who received a different treatment. Therefore, by weighting this patient by $1/\pi^{(A)}(X_i, \gamma)$ we can use their outcome μ to account for themselves as well as those $1/\pi^{(A)}(X_i, \gamma)-1$ other patients. Using this method to estimate the mean outcome conditional on

receiving A would result in the Inverse Probability of Treatment Weighting estimator [Horvitz and Thompson, 1952]

$$\hat{\mu}_{IPTW} = \frac{1}{n}\sum_{i=1}^{n}\frac{Z_i^{(A)}}{\pi^{(A)}(X_i,\hat{\gamma})}Y_i, \tag{2.1}$$

where $Z_i^{(A)}$ is the indicator for receiving treatment A, $Z_i^{(A)} = 1$, if the patient receives A, 0, otherwise, and Y_i is the outcome for the i^{th} subject, $i = 1, ..., n$. In the above equation we have also replaced γ by $\hat{\gamma}$ to indicate that when $\pi^{(A)}$ is unknown, it is estimated from the observed data.

Both of these estimating approaches provide a consistent estimator of the population mean μ given the correct specification of the models involved, outcome and treatment respectively. Misspecification of these models may result in biased and inefficient estimates. However, a third approach, doubly robust estimation, allows for consistent estimation as long as at least one of these models is correctly specified. This doubly robust estimator is given by

$$\hat{\mu}_{DR} = \frac{1}{n}\sum_{i=1}^{n}\left\{\frac{Z_i^{(A)}}{\pi^{(A)}(X_i,\hat{\gamma})}\,Y_i - \frac{Z_i^{(A)} - \pi^{(A)}(X_i,\hat{\gamma})}{\pi^{(A)}(X_i,\hat{\gamma})}m(X_i,\hat{\beta})\right\}. \tag{2.2}$$

For a formal argument for why this estimator is consistent when at least one of the two models is correctly specified, we refer readers to Robins Am. Statist. Assoc. Sect. Bayesian Statist. Sci [1999].

2.2.2 Causal Inference for Dynamic Treatment Regimes

Now consider the setting of two-stage dynamic treatment regimes where, at decision points, non-responders receive subsequent follow-up treatments and responders do

not. Patients will receive up to two levels of treatment. For simplicity, suppose there are only two first stage treatment options, A_1 and A_2, and two second stage treatment options, B_1 and B_2, for non-responders to the first stage treatments A_1 or A_2. We denote intermediate response with $R_i = 1$ denoting the i^{th} patient response and $R_i = 0$ for their non-response after the first stage of treatment. This setting allows four dynamic treatment regimes, namely, $d(A_j, B_k)$, $j, k = 1, 2$, where, for example, in regime $d(A_1, B_1)$, a patient initially receives treatment A_1 and if non-responsive is switched to treatment B_1. A patient who receives A_1, responds, and stays on that treatment would be considered to be on this regime, and so would someone that does not respond to A_1 and is switched to B_1. This creates an interesting problem when trying to estimate the mean regime outcome in even a randomized controlled trial. If there were another second stage treatment that a non-responder could switch to, B_2, then for patients who did not respond to A_1 and were switched to B_2, we do not observe the outcome under regime $d(A_1, B_1)$. However, they were part of that regime up until the second stage of treatment. This is equivalent to outcome data missing for these patients under regime $d(A_1, B_1)$ and hence when estimating the mean outcome under $d(A_1, B_1)$, we need to account for these missing data. If we simply take the average outcome of all patients who can be described as following regime $d(A_1, B_1)$ then we may get a biased estimate. In this average responders to A_1 will count for more than non-responders because some portion of non-responders to A_1, those assigned to B_2, have their outcome missing for our regime of interest.

Let us cast our problem in terms of counterfactual outcomes. Suppose $Y[d(A_j, B_k)]$ denotes the outcome had the patient been treated with the regime $d(A_j, B_k)$. In terms of this, the goal is to estimate $\mu[d(A_j, B_k)] = E\{Y[d(A_j, B_k)]\}$, the mean of a population who were treated with $d(A_j, B_k)$. However, $Y_i[d(A_j, B_k)]$ is not observed

for every individual, instead we only observe

$$(X_i, Z_i^{(A)}, R_i, (1-R_i)Z_i^{(B)}, Y_i; i = 1, ..., n) \ , \tag{2.3}$$

which are respectively the patient's observed covariates, first stage treatment assignment, response to first stage treatment, second stage treatment assignment for non-responders, and outcome.

The primary assumption required for causal inference is consistency (Robins et. al [2000]), which in a dynamic treatment regime setting means that the outcome Y_i a patient experiences under a regime $d(A_j, B_k)$ is in fact the counterfactual outcome $Y_i[d(A_j, B_k)]$. In the case of two treatments at each stage this would be

$$\begin{aligned} Y_i = Z_{1i}^{(A)}[R_i + (1-R_i) * Z_{1i}^{(B)}] * Y_i[d(A_1, B_1)] \\ +(1 - Z_{1i}^{(A)})[R_i + (1-R_i) * Z_{1i}^{(B)}] * Y_i[d(A_2, B_1)] \\ +Z_{1i}^{(A)}[R_i + (1-R_i) * (1 - Z_{1i}^{(B)})] * Y_i[d(A_1, B_2)] \\ +(1 - Z_{1i}^{(A)})[R_i + (1-R_i) * (1 - Z_{1i}^{(B)})] * Y_i[d(A_2, B_2)] \ . \end{aligned} \tag{2.4}$$

As previously mentioned, another assumption necessary for making causal inference is positivity. That is, that the probability of each treatment pathway being

experienced by a patient is non-zero. Staying in our two treatment per stage example this would mean

$$\begin{aligned}
P(Z_{1i}^{(A)} = 1) &> 0 \\
P(Z_{1i}^{(A)} = 0) &> 0 \\
P(R_i = 1 | Z_{1i}^{(A)} = 1) &> 0 \\
P(R_i = 0 | Z_{1i}^{(A)} = 1) &> 0 \\
P(R_i = 1 | Z_{1i}^{(A)} = 0) &> 0 \\
P(R_i = 0 | Z_{1i}^{(A)} = 0) &> 0 \\
P(Z_{1i}^{(B)} = 1 | Z_{1i}^{(A)} = 1, R_i = 0) &> 0 \\
P(Z_{1i}^{(B)} = 0 | Z_{1i}^{(A)} = 1, R_i = 0) &> 0 \\
P(Z_{1i}^{(B)} = 1 | Z_{1i}^{(A)} = 0, R_i = 0) &> 0 \\
P(Z_{1i}^{(B)} = 1 | Z_{1i}^{(A)} = 0, R_i = 0) &> 0 \ .
\end{aligned} \tag{2.5}$$

The last of our three assumptions is sequential randomization. This assumption states that the probability of being assigned to a treatment at a decision point depends only on information observed up to that time. More explicitly,

$$\begin{aligned}
P(Z_{ji}^{(A)} = 1) &= P(Z_{ji}^{(A)} = 1 | X_i) \\
P(Z_{ki}^{(B)} = 1) &= P(Z_{ki}^{(B)} = 1 | X_i, Z_{ji}^{(A)}).
\end{aligned} \tag{2.6}$$

Bembom and van der Laan [2006] describe a method, known as g-computation, originally described by Robins [1986], that uses weighted averages, based on response probability, of conditional intermediate responses to achieve unbiased estimates of

mean regime outcomes. Specifically, the mean outcome under regime $d(A_j, B_k)$ is estimated as

$$\hat{\mu}_G[d(A_j, B_k)] = \frac{1}{n}\sum_{i=1}^{n} \left\{ R_i m_j(X_i, \hat{\beta}_j) \ + (1 - R_i) m_{jk}(X_i, \hat{\beta_{jk}}) \right\}, \tag{2.7}$$

where R_i is the i^{th} patient's response after the first stage of treatment, $R_i = 1$, if the responded to first stage treatment, 0 otherwise, X_i is the i^{th} patient's covariates, X_i is the i^{th} patient's covariate information, X_i, $m_j(X_i, \beta_j)$ is the postulated outcome model for the conditional expectation of Y for responders given the covariates X_i, initial treatment A_j. Likewise, $m_{jk}(X_i, \beta_{jk})$ is the postulated outcome model for the conditional expectation of Y for non-responders given the covariates X_i, initial treatment A_j and second stage treatment B_k. We used hats to indicate that they are now estimated from the observed data.

The Inverse Probability of Treatment Weighted estimator in a two stage dynamic treatment regime setting becomes

$$\hat{\mu}_{IPTW}[d(A_j, B_k)] = \frac{1}{n}\sum_{i=1}^{n} \left\{ \frac{Z_{ji}^{(A)}}{\pi_j^{(A)}(X_i, \hat{\gamma_j})} \left[R_i + \frac{(1 - R_i) Z_{ki}^{(B)}}{\pi_{jk}^{(B)}(X_i, \hat{\gamma_{jk}})} \right] \right\} Y_i \ , \tag{2.8}$$

where $\pi_j^{(A)}(X_i, \hat{\gamma_j})$ is the modeled probability of receiving treatment A_j based on the covariates X_i, and $\pi_{jk}^{(B)}(X_i, \hat{\gamma_{jk}})$ is the modeled probability of receiving treatment B_k based on the covariates X_i given non-response to $Z_{ji}^{(A)}$.

Lastly consider a locally efficient doubly robust estimator for a two-stage dynamic treatment regime introduced by Wahed and Tsiatis [2004].

$$\begin{aligned}\hat{\mu}_{DR}[d(A_j, B_k)] = \sum_{i=1}^{n} Z_{ji}^{(A)} \left\{ \left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\pi_{jk}^{(B)}(X_i, \hat{\gamma_{jk}})} \right] Y_i \right. \\ \left. - \frac{Z_{ki}^{(B)} - \pi_{jk}^{(B)}(X_i, \hat{\gamma_{jk}})}{\pi_{jk}^{(B)}(X_i, \hat{\gamma_{jk}})} (1-R_i) m_{jk}(X_i, \beta_{jk}) \right\} / \sum_{i=1}^{n} Z_{ji}^{(A)} \ .\end{aligned} \tag{2.9}$$

This estimator is consistent provided at least one of the two models π_{jk} and m_{jk} are correctly specified. When both models are correctly specified, this estimator is most efficient, provided β_j and β_{jk} are consistently estimated. When we expand this estimator to include an inverse probability of treatment weighting component for first stage of treatment, A_j this estimator becomes.

$$\begin{aligned}\hat{\mu}_{DR}[d(A_j, B_k)] = \frac{1}{n} \sum_{i=1}^{n} \left\{ \frac{Z_{ji}^{(A)}}{\hat{\pi}_j^{(A)}} \left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\hat{\pi}_{jk}^{(B)}} \right] Y_i \right. \\ \left. - \frac{Z_{ji}^{(A)} - \hat{\pi}_j^{(A)}}{\hat{\pi}_j^{(A)}} R_i m_j(X_i, \hat{\beta}_j) \ - \frac{Z_{ji}^{(A)} Z_{ki}^{(B)} - \hat{\pi}_j^{(A)} \hat{\pi}_{jk}^{(B)}}{\hat{\pi}_j^{(A)} \hat{\pi}_{jk}^{(B)}} (1-R_i) m_{jk}(X_i, \hat{\beta}_{jk}) \right\} ,\end{aligned} \tag{2.10}$$

where $m_j(X_i, \hat{\beta}_j)$ and $m_{jk}(X_i, \hat{\beta}_{jk})$ are the estimated mean outcomes under treatment pathways A_j and $A_j B_k$ respectively.

This estimator is biased and inefficient. It may require that first stage treatment assignment, $Z_{ji}^{(A)}$ is independent of response to that treatment, R_i and violating this assumption might be to blame for the problem with bias and inefficiency.

2.3 PROPOSED ESTIMATORS

2.3.1 The Hybrid Inverse Probability of Treatment Weighted Estimator

From the causal inference point of view, we are interested in the outcome if the study population had received regime $d(A_j, B_k)$. As discussed earlier, this quantity is $Y_i[d(A_j, B_k)]$. Since all individuals in the sample did not follow this particular regime, $Y_i[d(A_j, B_k)]$ is unobserved for individuals whose treatment is inconsistent with this regime, and hence missing data techniques can be applied to estimate the parameter of interest, $\mu = \mu_{Y_i[d(A_j,B_k)]} = E[Y_i[d(A_j, B_k)]]$. We first define R^*_{jki}, an indicator that takes the value 1 when the i^{th} individual in the sample is consistent with the treatment regime $d(A_j, B_k)$, and 0 otherwise. More specifically,

$$R^*_{jki} = Z^{(A)}_{ji}[R_i + (1 - R_i) * Z^{(B)}_{ki}] \ . \tag{2.11}$$

The probability of being on the regime $d(A_j, B_k)$, $P(R^*_{jki} = 1)$ similarly expressed as a composite of the probabilities for treatment assignment and intermediate response,

$$P(R^*_{jki} = 1) = P(Z^{(A)}_{ji}[R_i + (1 - R_i) * Z^{(B)}_{ki}] = 1) =$$
$$P(Z^{(A)}_{ji} = 1)[P(R_i = 1|Z^{(A)}_{ji} = 1)+(1 - P(R_i = 1|Z^{(A)}_{ji} = 1)) * P(Z^{(B)}_{ki} = 1|R_i = 1, Z^{(A)}_{ji} = 1)] \ .$$

If we were to utilize the probability of being on the regime of interest as a unique weight for each patient we could then define this quantity as

$$\pi^{(R^*)}_{jk}(X_i, \hat{\gamma_{jk}}) = \pi^{(A)}_{j}(X_i, \hat{\gamma_{A_j}})[\pi^{(R)}_{j|Z^{(A)}_{ji}=1}(X_i, \hat{\gamma_R})+(1-\pi^{(R)}_{j|Z^{(A)}_{ji}=1}(X_i, \hat{\gamma_R}))*\pi^{(B)}_{k|Z^{(A)}_{ji}=1,R=1}(X_i, \hat{\gamma_{B_{jk}}})] \ . \tag{2.12}$$

The terms $\pi_j^{(A)}(X_i, \hat{\gamma_{A_j}})$, $\pi^{(B)}_{k|Z_{ji}^{(A)}=1,R=1}(X_i, \gamma_{\hat{B}_{jk}})$, and $\pi^{(R)}_{j|Z_{ji}^{(A)}=1}(X_i, \hat{\gamma_R})$ can in turn be estimated using logistic regression or similar models. An Inverse Probability of Treatment Weighted estimator can then be defined as

$$\hat{\mu}_{HIPTW}[d(A_j, B_k)] = \frac{1}{n}\sum_{i=1}^{n}\left\{ \frac{R^*_{jki}}{\hat{\pi}^{(R^*)}_{jk}(X_i, \hat{\gamma_{jk}})} \right\} Y_i \ . \tag{2.13}$$

We will call this new estimator the Hybrid Inverse Probability of Treatment Weighted estimator (HIPTW). We call it a hybrid estimator because now we are modeling both treatment and intermediate response.

2.3.2 The Hybrid Doubly Robust Estimator

The extension of the HIPTW estimator, (2.13), to a doubly robust variant is straightforward and analogous to the extension in the single stage setting. We have our Inverse Probability of Treatment Weighted term and we augment it with a model based estimation term that only comes in to play when the probability of being on the regime of interest is incorrectly modeled. We call this the Hybrid Doubly Robust Estimator and define it as

$$\hat{\mu}_{HDR}[d(A_j, B_k)] = \frac{1}{n}\sum_{i=1}^{n}\left\{ \frac{R^*_{jki}}{\hat{\pi}^{(R^*)}_{jk}(X_i, \hat{\gamma_{jk}})} Y_i - \frac{R^*_{jki} - \hat{\pi}^{(R^*)}_{jk}(X_i, \hat{\gamma_{jk}})}{\hat{\pi}^{(R^*)}_{jk}(X_i, \hat{\gamma_{jk}})} \hat{\mu}^{R^*}_{jk}(X_i, \hat{\beta}) \right\}, \tag{2.14}$$

where $\mu^{R^*}_{jk}(X_i, \beta)$ is the mean of the counterfactual outcome corresponding to being treated with the regime $d(A_j, B_k)$. This new, hybrid two stage doubly robust estimator has the advantage of being more efficient than the existing two-stage doubly robust estimator. Additionally it does not require the previously discussed

unreasonable assumption of first stage treatment assignment is independent of the outcome of that treatment. Additionally (2.14) can be viewed as the solution to the estimating equations

$$\psi_i(\theta) = \begin{pmatrix} \psi_{jiA}(X_i, Z_{ji}^{(A)}; \gamma_{A_j}) \\ \psi_{jkiB}(X_i, Z_{ji}^{(A)}, Z_{ki}^{(B)}; \gamma_{B_{jk}}) \\ \psi_{iR}(X_i, Z_{ji}^{(A)}, R_i; \gamma_R) \\ \psi_{iR^*}(X_i, R_{jki}^*, Y_i; \beta) \\ \psi_{iHDR}(X_i, R_{jki}^*, Y_i; \mu_{HDR}) \end{pmatrix} , \tag{2.15}$$

where the ψ functions inside the parenthesis are individual estimating equations for the components of the combined parameter vector,

$$\theta = (\gamma_{A_j}, \gamma_{B_{jk}}, \gamma_R, \beta, \mu_{HDR}) , \tag{2.16}$$

used in the estimation of $\mu_{HDR}[d(A_j, B_k)]$. The first three correspond to logistic regression models while the fourth is the estimating equation for a linear regression model for complete cases. More explicitly,

$$\psi_{jiA}(X_i, Z_{ji}^{(A)}; \gamma_{A_j}) = X_i \left(Z_{ji}^{(A)} - \frac{1}{1 + e^{-X_i^T \gamma_{A_j}}} \right)$$

$$\psi_{jkiB}(X_i, Z_{ji}^{(A)}, Z_{ki}^{(B)}; \gamma_{B_{jk}}) = X_i \left(Z_{jki}^{(B)} - \frac{1}{1 + e^{-(X_i^T, Z_{ji}^{(A)})\gamma_{B_{jk}}}} \right)$$

$$\psi_{iR}(X_i, Z_{ji}^{(A)}, R_i; \gamma_R) = X_i \left(R_i - \frac{1}{1 + e^{-(X_i^T, Z_{ji}^{(A)})\gamma_R}} \right)$$

$$\psi_{iR^*}(X_i, R_{jki}^*, Y_i; \beta) = R_{jki}^* (Y_i - \mu_{jk}^{R^*}(X_i, \hat{\beta}))$$

$$\psi_{iHDR}(X_i, R_{jki}^*, Y_i; \mu_{HDR}) = \frac{R_{jki}^*}{\pi_{jk}^{(R^*)}(X_i, \gamma_{jk})} Y_i - \frac{R_{jki}^* - \pi_{jk}^{(R^*)}(X_i, \gamma_{jk})}{\pi_{jk}^{(R^*)}(X_i, \gamma_{jk})} \mu_{jk}^{R^*}(X_i, \beta) - \mu_{HDR}[d(A_j, B_k)] .$$

The final equation is the estimating equation corresponding to the Hybrid Doubly Robust estimator itself.

Defining the Hybrid Doubly Robust estimator as a solution to the above set of estimating equations allows us to classify it as an M-estimator. It follows then that it is asymptotically normal, as well as consistent (Stefanski and Boos [2002]). An estimator of the variance of the Hybrid Double Robust estimator can be found with the sandwich estimator of the variance $V_n = A_n(\hat{\theta})^{-1} B_n(\hat{\theta}) A_n(\hat{\theta})^{-1^T}/n$, where

$$A_n(\theta) = \frac{1}{n}\sum_{i=1}^{n} \frac{\partial \psi_i(\theta)}{\partial \theta} = \frac{1}{n}\sum_{i=1}^{n} \begin{pmatrix} \frac{\partial \psi_{jiA}}{\partial \gamma_{A_j}} & 0 & 0 & 0 & 0 \\ 0 & \frac{\partial \psi_{jkiB}}{\partial \gamma_{B_j k}} & 0 & 0 & 0 \\ 0 & 0 & \frac{\partial \psi_{iR}}{\partial \gamma_R} & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ \frac{\partial \psi_{iHDR}}{\partial \gamma_{A_j}} & \frac{\partial \psi_{iHDR}}{\partial \gamma_{B_{jk}}} & \frac{\partial \psi_{iHDR}}{\partial \gamma_R} & \frac{\partial \psi_{iHDR}}{\partial \beta} & 1 \end{pmatrix}$$

and

$$B_n(\theta) = \frac{1}{n}\sum_{i=1}^{n} \psi_i(\theta)\psi_i(\theta)^T,$$

where the derivatives found in the matrix $A_n(\theta)$ are defined as

$$\frac{\partial \psi_{jiA}}{\partial \gamma_{A_j}} = X_i e^{-X_i^T \gamma_{A_j}} \left(\frac{1}{1 + e^{-X_i^T \gamma_{A_j}}} \right)^2 X_i^T$$

$$\frac{\partial \psi_{jkiB}}{\partial \gamma_{B_{jk}}} = X_i e^{-X_i^T \gamma_{B_{jk}}} \left(\frac{1}{1 + e^{-X_i^T \gamma_{B_k}}} \right)^2 X_i^T$$

$$\frac{\partial \psi_{iR}}{\partial \gamma_R} = X_i e^{-X_i^T \gamma_R} \left(\frac{1}{1 + e^{-X_i^T \gamma_R}} \right)^2 X_i^T$$

$$\frac{\partial \psi_{iHDR}}{\partial \gamma_{A_j}} = \frac{\pi_{\gamma Aj}^{R*} R_i * (Y_i - \mu_{jk}^R *)}{(\pi^{R^*})^2}$$

$$\frac{\partial \psi_{iHDR}}{\partial \gamma_{B_{jk}}} = \frac{\pi_{\gamma Bk}^{R*} R_i * (Y_i - \mu_{jk}^R *)}{(\pi^{R^*})^2}$$

$$\frac{\partial \psi_{iHDR}}{\partial \gamma_R} = \frac{\pi_{\gamma R}^{R*} R_i * (Y_i - \mu_{jk}^R *)}{(\pi^{R^*})^2}$$

$$\frac{\partial \psi_{iHDR}}{\partial \beta} = \frac{R_i - \pi_{R^*}}{\pi_{R^*}}$$

$$\pi_{\gamma Aj}^{R*} = \frac{\partial \pi_{jk}^{(R^*)}(X_i, \gamma_{jk})}{\partial \gamma_{A_j}}$$

$$\pi_{\gamma Bjk}^{R*} = \frac{\partial \pi_{jk}^{(R^*)}(X_i, \gamma_{jk})}{\partial \gamma_{B_j k}}$$

$$\pi_{\gamma R}^{R*} = \frac{\partial \pi_{jk}^{(R^*)}(X_i, \gamma_{jk})}{\partial \gamma_R} \ .$$

A similar approach can be used to find an estimator of the variance for the Hybrid Inverse Probability of Treatment Weighted estimator, the difference being that the $A_n(\theta)$ and $B_n(\theta)$ matrices have smaller dimensions.

2.4 SIMULATIONS

We based our simulation on the one found in Kang and Schafer (2007) with a few notable exceptions. For each subject we generate 8 independent identically distributed standard normal random variables, which we will call X_a, X_b, X_c, X_d, X_e, X_f, X_g, and X_h. We assume a two stage design, with non-responders from the first stage being rerandomized to the second stage. Without loss of generality we will be dropping the j and k subscript, concerning ourselves only with the regime where $Z_A = 1$ and $Z_B = 1$. The treatment probabilities for each stage are

$$\pi_A = expit(\gamma_{10} + \gamma_{11} * X_e + \gamma_{12} * X_f + \gamma_{13} * X_g + \gamma_{14} * X_h) \tag{2.17}$$

and

$$\pi_B = expit(\gamma_{20} + \gamma_{21} * X_e + \gamma_{22} * X_f + \gamma_{23} * X_g + \gamma_{24} * X_h + \gamma_{25} * Z_A) \ , \tag{2.18}$$

where *expit* is the inverse logit function.

With probability of response to the first stage being defined as

$$\pi_R \sim BERNOULLI(\tau_0 + \tau_1 * Z_A) \ . \tag{2.19}$$

The outcome variable, Y, is then written as

$$\begin{aligned} Y = \beta_0 + \beta_1 * Z_A + \beta_2 * Z_B + \beta_3 * X_a + \beta_4 * X_b \\ +\beta_5 * X_c + \beta_6 * X_d + \beta_7 * R + \beta_8 * Z_A * Z_B + e \ , \end{aligned} \tag{2.20}$$

where e $\sim N(0,1)$ are normally distributed error terms. Now the observed covariates are

$$\begin{aligned} X_1 &= e^{\frac{X_a}{2}} \\ X_2 &= \frac{X_b}{1+e^{X_a}} + 10 \\ X_3 &= (\frac{X_a \times X_c}{25} + .6)^3 \\ X_4 &= (X_b + X_d + 20)^2 \\ X_5 &= e^{\frac{X_e}{2}} \\ X_6 &= \frac{X_f}{1+e^{X_e}} + 10 \\ X_7 &= (\frac{X_e \times X_g}{25} + .6)^3 \\ X_8 &= (X_f + X_h + 20)^2 \ . \end{aligned} \tag{2.21}$$

Since it would be unreasonably difficult for any investigator to correctly identify the correct relationship between observed covariates and either outcome or treatment we assume the following incorrect models could be specified:

$$Y = \hat{\xi_{10}} + \hat{\xi_{11}} * X_1 + \hat{\xi_{12}} * X_2 + \hat{\xi_{13}} * X_3 + \hat{\xi_{14}} * X_4 + \hat{\xi_{15}} * Z_A \ , \tag{2.22}$$

which is the outcome model for those who did respond to the initial treatment.

$$Y = \hat{\xi_{20}} + \hat{\xi_{21}} * X_1 + \hat{\xi_{22}} * X_2 + \hat{\xi_{23}} * X_3 + \hat{\xi_{24}} * X_4 + \hat{\xi_{25}} * Z_A + \hat{\xi_{26}} * Z_B + \hat{\xi_{27}} * Z_A * Z_B \ , \tag{2.23}$$

which is the outcome model for those who did not respond to the initial treatment and were rerandomized to a second stage treatment.

$$\begin{aligned} Z_A &\sim \hat{\eta_{11}} * X_5 + \hat{\eta_{12}} * X_6 + \hat{\eta_{13}} * X_7 + \hat{\eta_{14}} * X_8 \\ &\textit{and} \\ Z_B &\sim \hat{\eta_{21}} * X_5 + \hat{\eta_{22}} * X_6 + \hat{\eta_{23}} * X_7 + \hat{\eta_{24}} * X_8 + \hat{\eta_{25}} * Z_A \ , \end{aligned} \tag{2.24}$$

which are logistic regressions of the treatment assignments on the covariates with ξ and η as their estimated coefficients.

We perform a series of 10,000 Monte Carlo simulations in order to asses the validity of our estimator under circumstances where one or both of the types of models are misspecified. For the first simulation the true value of the estimand, the mean counterfactual outcome $\mu\{Y[d(A,B)]\} = E\{Y[d(A,B)]\} = 142.398$,
$\beta = (110, 24.7, 13.7, 124.7, 73.7, 23.7, 50.7, 20.2, -10.7)$, $\tau = (.2, -.7)$, and
$\gamma = (\gamma_{10}, \gamma_{11}, \gamma_{12}, \gamma_{13}, \gamma_{14}, \gamma_{20}, \gamma_{21}, \gamma_{22}, \gamma_{23}, \gamma_{24}, \gamma_{25}) = (-.3, -.75, -.2, .5, 1, .2, -.2, -.3, .4, -.5, -.7)$.
In the second simulation, created to more closely mimic the STAR*D data, the true value of the estimand is $\mu\{Y[d(A,B)]\} = E\{Y[d(A,B)]\} = 7.39546$,
$\beta = (-0.932, -0.137, 1.51, 0.045, 0.313, 0, 0, 5.09, -0.951)$, $\tau = (-0.559, 0.711)$, and
$\gamma = (0, -0.457, -0.135, 0, 0, 0, -0.074, -0.041, 0, 0, .500)$. We also varied sample size in both simulations, running the simulation at $n = 200$ and $n = 500$ for the first simulation and n=200 for the second, as well as examining the estimators under both standard normal and t-distribution error structures in the first simulation and standard normal error structures in the second.

2.4.1 Simulation Results

In Table 1 we see the results for running the simulation at $n = 200$ with standard normal errors. When both treatment and outcome models are correct, all the estimators except for the existing doubly robust estimator are approximately unbiased with relative biases below 3%. For the existing doubly robust estimator, this bias is 4.2% vesus a bias of -2.6% for the Hybrid Doubly Robust estimator. The G-Computation and IPTW estimators expectedly have the smaller biases compared to others, 0.4 and -0.5% respectively, while the Hybrid IPTW estimator has a larger bias at -2.9%.

Table 1: n=200, Standard Normal Errors, $\mu\{Y[d(A,B)]\} = 142.398$

Treatment Model Outcome Model	Method	Bias (%)	MCSE	MSE
	$G-Comp$	0.4	21.4	478
Correct	$IPTW$	-0.5	38.1	1490
Correct	DR	4.2	20.2	463
	$HIPTW$	-2.9	31.5	1040
	HDR	-2.6	11.2	150
	$G-Comp$	0.4	21.4	478
Incorrect	$IPTW$	4.4	63.6	4140
Correct	DR	0.9	20.1	423
	$HIPTW$	2.1	66.6	4510
	HDR	-2.5	11.3	151
	$G-Comp$	0.4	21.4	478
Correct	$IPTW$	0.5	38.1	1490
Incorrect	DR	4.2	20.2	463
	$HIPTW$	-2.9	31.5	1040
	HDR	-1.6	13.7	207
	$G-Comp$	0.4	21.4	478
Incorrect	$IPTW$	4.4	63.6	4140
Incorrect	DR	3.3	36.9	1420
	$HIPTW$	2.1	66.6	4510
	HDR	-1.6	17.1	315

$G-Comp$ = G-Computation, $IPTW$ = Inverse Probability of Treatment Weighted Estimator, DR = Doubly Robust Estimator, $HIPTW$ = Hybrid Inverse Probability of Treatment Weighted Estimator, HDR = Hybrid Doubly Robust Estimator

Table 2: n=500, Standard Normal Errors, $\mu\{Y[d(A,B)]\} = 142.398$

Treatment Model Outcome Model	Method	Bias (%)	MCSE	MSE
	$G-Comp$	0.7	13.3	191
Correct	$IPTW$	0.6	24.6	628.0
Correct	DR	3.0	12.5	186.0
	$HIPTW$	-2.4	19.2	399.0
	HDR	-2.4	7.1	68.1
	$G-Comp$	0.7	13.3	191
Incorrect	$IPTW$	8.8	144.0	21100.0
Correct	DR	-1.9	15.7	270.0
	$HIPTW$	6.1	156.0	24600.0
	HDR	-2.3	8.1	83.5
	$G-Comp$	0.7	13.30	191
Correct	$IPTW$	0.6	24.6	628.0
Incorrect	DR	3.0	12.5	186.0
	$HIPTW$	-2.4	19.2	399.0
	HDR	-1.9	8.5	88.8
	$G-Comp$	0.7	13.3	191
Incorrect	$IPTW$	8.8	144.0	21100
Incorrect	DR	-0.5	17.4	321
	$HIPTW$	6.1	156.0	24600
	HDR	-1.5	33.9	1180

$G-Comp$ = G-Computation, $IPTW$ = Inverse Probability of Treatment Weighted Estimator, DR = Doubly Robust Estimator, $HIPTW$ = Hybrid Inverse Probability of Treatment Weighted Estimator, HDR = Hybrid Doubly Robust Estimator

Table 3: n=200, t-Distribution with df=5 Errors, $\mu\{Y[d(A,B)]\} = 142.398$

Treatment Model Outcome Model	Method	Bias (%)	MCSE	MSE
	$G-Comp$	0.5	21.4	479
Correct	$IPTW$	-0.1	46.4	2200
Correct	DR	4.6	37.6	1500
	$HIPTW$	-3.0	31.5	1040
	HDR	-2.4	11.1	147
	$G-Comp$	0.5	21.4	479
Incorrect	$IPTW$	5.3	136.0	18600
Correct	DR	0.6	37.2	1420
	$HIPTW$	2.5	75.5	5780
	HDR	-2.3	11.8	163
	$G-Comp$	0.5	21.4	479
Correct	$IPTW$	-0.1	46.4	2200
Incorrect	DR	4.6	37.6	1500
	$HIPTW$	-3.0	31.5	1040
	HDR	-1.5	13.7	205
	$G-Comp$	0.5	21.4	479
Incorrect	$IPTW$	5.3	136.0	18600
Incorrect	DR	3.7	50.9	2670
	$HIPTW$	2.5	75.5	5780
	HDR	-1.4	24.1	607

$G-Comp$ = G-Computation, $IPTW$ = Inverse Probability of Treatment Weighted Estimator, DR = Doubly Robust Estimator, $HIPTW$ = Hybrid Inverse Probability of Treatment Weighted Estimator, HDR = Hybrid Doubly Robust Estimator

Table 4: n=500, t-Distribution with df=5 Errors, $\mu\{Y[d(A,B)]\} = 142.398$

Treatment Model Outcome Model	Method	Bias (%)	MCSE	MSE
	$G-Comp$	0.7	13.4	194
Correct	$IPTW$	0.311	23.5	576.0
Correct	DR	3.0	12.4	184.0
	$HIPTW$	-2.6	18.8	385.0
	HDR	-2.4	7.0	67.9
	$G-Comp$	0.7	13.4	194
Incorrect	$IPTW$	-9.5	2390.0	5690000
Correct	DR	-1.7	15.3	256
	$HIPTW$	-17.1	3050.0	9280000
	HDR	-2.9	86.7	7610
	$G-Comp$	0.7	13.4	194
Correct	$IPTW$	0.3	23.5	576.0
Incorrect	DR	3.0	12.4	184.0
	$HIPTW$	-2.6	18.8	385.0
	HDR	-2.0	8.5	89.6
	$G-Comp$	0.7	13.4	194
Incorrect	$IPTW$	-9.5	2390.0	5690000
Incorrect	DR	-0.3	16.8	300
	$HIPTW$	-17.1	3050.0	9280000
	HDR	-6.1	620.0	385000

$G-Comp$ = G-Computation, $IPTW$ = Inverse Probability of Treatment Weighted Estimator, DR = Doubly Robust Estimator, $HIPTW$ = Hybrid Inverse Probability of Treatment Weighted Estimator, HDR = Hybrid Doubly Robust Estimator

Table 5: With X_8, n=200, Standard Normal Errors, $Y[d(A,B)] = 142.398$

Treatment Model Outcome Model	Method	Sandwich Est. SE	MCSE	Coverage (%)
Correct Correct	HDR	11.1	11.2	93.7
Incorrect Correct	HDR	128.1	11.3	99.8
Correct Incorrect	HDR	13.5	13.7	95.8
Incorrect Incorrect	HDR	610.2	17.1	100

HDR = Hybrid Doubly Robust Estimator

Table 6: With X_8, n=500, Standard Normal Errors, $Y[d(A,B)] = 142.398$

Treatment Model Outcome Model	Method	Sandwich Est. SE	MCSE	Coverage (%)
Correct Correct	HDR	7.0	7.1	92.0
Incorrect Correct	HDR	63.5	8.1	99.8
Correct Incorrect	HDR	8.5	8.5	93.1
Incorrect Incorrect	HDR	310.0	33.9	100

HDR = Hybrid Doubly Robust Estimator

Table 7: Without X_8, n=200, Standard Normal Errors, $Y[d(A,B)] = 142.398$

Treatment Model Outcome Model	Method	Sandwich Est. SE	MCSE	Coverage (%)
Correct Correct	HDR	11.1	11.2	93.7
Incorrect Correct	HDR	12.0	11.2	95.5
Correct Incorrect	HDR	13.5	13.7	95.8
Incorrect Incorrect	HDR	20.5	14.2	98.0

HDR = Hybrid Doubly Robust Estimator

Table 8: Without X_8, n=500, Standard Normal Errors, $Y[d(A,B)] = 142.398$

Treatment Model Outcome Model	Method	Sandwich Est. SE	MCSE	Coverage (%)
Correct Correct	HDR	7.0	7.1	92.5
Incorrect Correct	HDR	7.4	8.1	93.4
Correct Incorrect	HDR	8.5	8.5	93.1
Incorrect Incorrect	HDR	10.8	9.19	96.9

HDR = Hybrid Doubly Robust Estimator

Table 9: n=2000, Standard Normal Errors, $\mu\{Y[d(A,B)]\} = 142.398$

Treatment Model	Method	Bias	MCSE	MSE
Outcome Model				
	$G-Comp$	-0.002	0.135	0.153
Correct	$IPTW$	0.007	0.146	0.167
Correct	DR	9.78	0.116	0.653
	$HIPTW$	-8.00	0.158	0.533
	HDR	-8.04	0.150	0.526
	$G-Comp$	-0.002	0.135	0.153
Incorrect	$IPTW$	-0.06	0.158	0.183
Correct	DR	8.10	0.157	0.541
	$HIPTW$	-6.56	0.191	0.463
	HDR	-8.04	0.161	0.540
	$G-Comp$	-0.002	0.135	0.153
Correct	$IPTW$	0.007	0.146	0.167
Incorrect	DR	9.78	0.116	0.653
	$HIPTW$	-8.00	0.158	0.533
	HDR	-8.03	0.151	0.527
	$G-Comp$	-0.00213	0.135	0.153
Incorrect	$IPTW$	-0.006	0.158	0.183
Incorrect	DR	8.10	0.157	0.540
	$HIPTW$	-6.56	0.191	0.463
	HDR	-8.04	0.161	0.541

$G-Comp$ = G-Computation, $IPTW$ = Inverse Probability of Treatment Weighted Estimator, DR = Doubly Robust Estimator, $HIPTW$ = Hybrid Inverse Probability of Treatment Weighted Estimator, HDR = Hybrid Doubly Robust Estimator

Table 10: n=2000, Standard Normal Errors, $Y[d(A,B)] = 7.39546$

Treatment Model Outcome Model	Method	Sandwich Est. SE	MCSE	Coverage (%)
Correct Correct	HDR	0.146	0.150	85.0
Incorrect Correct	HDR	0.138	0.161	80.1
Correct Incorrect	HDR	0.146	0.151	83.8
Incorrect Incorrect	HDR	0.156	0.161	85.8

HDR = Hybrid Doubly Robust Estimator

Likewise, all of the estimators except for the IPTW and Hybrid IPTW have Monte Carlo standard errors under 30. The existing IPTW estimator has a MCSE of 38.1 compared to the Hybrid IPTW's MCSE of 31.5. The Hybrid Doubly Robust estimator performs the best with a MCSE of 11.2 versus 20.2 for the existing Doubly Robust estimator and 21.4 for the G-Computation estimator. As in the case of the MCSE, the MSE of the existing IPTW and Hybrid IPTW estimators is much larger than the other estimators, at 1490 and 1040 respectively, while the other estimators have an MSE under 500. The Hybrid Doubly Robust estimator, with an MSE of 150, also outperforms the existing Doubly Robust and G-Computation estimators, whose MSEs are 463 and 478, respectively. When only the outcome models are incorrectly specified the estimators maintain their relationships to one another in terms of bias, Monte Carlo standard error, and mean squared error.

However, we see some of these relationships change when only the treatment models are incorrectly specified. The existing Doubly Robust estimator now has a lower bias than the Hybrid Doubly Robust estimator at 0.9% and -2.5% respectively. The Hybrid IPTW estimator also now has lower bias, at 2.1%, than both the Hybrid Doubly Robust estimator and the existing IPTW estimator, which has a bias of 4.4%. Also, now the relationship between the existing IPTW and Hybrid IPTW estimators is reversed in terms of both MCSE and MSE. The existing IPTW estimator now has a lower MCSE and MSE, 63.6 and 4510 respectively. While the Hybrid IPTW estimator has an MCSE of 66.6 and an MSE of 4510. The other estimators keep their relationships to one another and maintain Monte Carlo standard errors under 30 and mean squared errors under 500.

When both treatment and outcome models are incorrectly specified the relationships between these estimators change slightly from when only the treatment models were incorrectly specified. In terms of relative bias the existing IPTW and Doubly

Robust estimators are now both over 3% with 4.4% and 3.1% respectively. The Hybrid IPTW estimator is now less biased, at 2.1%, than both the existing IPTW and Doubly Robust estimators. Additionally the existing Doubly Robust estimator now has an MCSE over 30, at 36.9, and an MSE over 500, at 1420. This means that the G-Computation estimator now has a lower MSE, at 478, than the existing Doubly Robust estimator in addition to the IPTW and Hybrid IPTW estimators. Now only the Hybrid Doubly Robust estimator has a lower MSE, at 315, than the G-Computation estimator.

In summary, Table 1 shows us some interesting dynamics between the estimators. When the treatment models are correctly specified, the regular IPTW estimator is less biased, has a higher Monte Carlo standard error, and a lower MSE than the new Hybrid IPTW estimator. This is result, however is reversed when the treatment model is misspecified. The Hybrid Doubly Robust estimator has the lowest Monte Carlo standard errors and MSE of any estimator regardless of model misspecification at this sample size. The regular Doubly Robust estimator has lower bias than the Hybrid Doubly Robust estimator only when the treatment models are incorrectly specified and the outcome models are not. G-computation has the lowest bias of any of the methods but it is important to note that only the outcome models for pathway endpoints are incorrectly specified, not the models for intermediate response. G-computation and regular IPTW produce estimators with biases of similar magnitudes when both treatment and outcome models are correctly specified.

Table 2, where now $n = 500$ with standard normal errors has similar results to Table 1 with a few notable exemptions. Now when both the oucome and treatment models are incorrectly specified the Hybrid Doubly Robust estimator performs worse than the regular Doubly Robust estimator in terms of bias (-1.5% versus -.05%), MCSE (33.9 versus 17.4), and MSE (1180 vs 321).

Table 3 are the results for running the simulation at $n = 200$ with error terms being generated from a student's t-distribution with 5 degrees of freedom. It is similar to Table 1 with the exception that now when both outcome and treatment models are incorrectly specified the Hybrid Doubly Robust estimator is inferior to the G-computation estimator in terms of MCSE (24.1 versus 21.4) and MSE (607 versus 479).

Table 4 produces interesting results. One of the Monte Carlo simulation runs contained a single observation that, under the incorrectly specified model for estimating probability of receiving treatment A, received a very small probability of receiving its actual treatment assignment. This in turn caused that observation to receive a abnormally high weight. The estimators utilizing inverse probability weighting for this run suffered from increased estimated biases, Monte Carlo standard errors, and MSEs as a result of this single run.

Table 5 shows a comparison between the stadard error computed using the sandwich estimator of the variance and the Monte Carlo standard error at $n = 200$. The Table shows that when the treatment models are correctly specified the sandwich estimator produces accurate, if slightly under estimated, estimators of the standard error. When both models are correctly specified the sandwich estimator is 11.1 versus the MCSE of 11.2 and when only the outcome is incorrectly specified the sandwich estimator is 13.5 versus the MCSE of 13.7. When the treatment models are incorrectly specified then the sandwich estimator of the variance appears not to be very accurate at all. With only the treatment models incorrectly specified the sandwich estimator is 128.1 versus the MCSE of 11.3. When both models are incorrectly specified the sandwich estimator becomes even less accurate with 610.2 versus the MCSE of 17.1. Table 6 corroborates this by showing similar discrepancies under treatment model misspecification at $n = 500$. The coverage at both sample sizes are provides slightly

suboptimal coverage when the treatment model is correctly specified and over coverage when incorrect. A close examination of what was happening in the simulation revealed that the elements of the matrix $A_n(\hat{\theta})$ corresponding to X_8, a variable in the misspecified treatment models, were abnormally large and were contributing the majority of the estimated error.

Tables 7 and 8 show the results of the simulation with X_8 removed from the misspecified treatment models at $n = 200$ and $n = 500$ respectively. When the treatment models are incorrectly specified and the outcome models are not the sandwich estimator is accurate, with it being a slight over estimate at $n = 200$, with a value of 12.0 versus the MCSE of 11.2, and a slight under estimate at $n = 500$, with a value of 7.4 versus the MCSE of 8.1. Both Tables show that when both the treatment models and the outcome models are incorrectly specified the standard error derived from the sandwich estimator of the variance are slight over estimates of the standard error. This over estimate is higher at the smaller sample size with a sandwich estimator value of 20.5 versus the MCSE of 14.2. At $n = 500$ this over estimated value is lower at 10.8 versus 9.19. At $n = 200$ the coverage probabilities when at least one model is correctly specified are close to 95% while at $n = 200$ they become slightly suboptimal. When both models are misspecified the coverage at both sample sizes are larger then 95%.

Finally, in Tables 9 and 10 we see that in the second simulation the HDR estimator performs worse then the other estimators. There is an increase in bias, likely caused by the smaller ratio of intercept to treatment effect coefficients in the outcome model. The higher effect of first stage treatment on outcome reveals a problem with the Hybrid estimators in such a scenario. The lower coverage for HDR in the the second simulation is due to the higher bias of the estimator.

2.5 CONCLUSION

At low sample sizes, a number of encouraging results start to appear. First, the new hybrid estimators are more efficient than their non-hybrid counterparts. Also, at low sample sizes, and with the exception of when only the outcome is correctly specified, the new Hybrid Double Robust estimator is both less biased and has a smaller mean squared error then the classic doubly robust estimator. The Hybrid Inverse Probability of Treatment Weighted estimator is more biased than its non-hybrid counterpart under correct model specification but considerably less biased when the treatment model is incorrectly specified. The hybrid inverse probability weighted estimator has a lower mean squared than its non-hybrid counterpart regardless of correct treatment model specification.

At larger sample sizes, the Hybrid Double Robust estimator is less biased, more efficient, and has a lower mean squared error than the non-Hybrid Double Robust estimator but only when the treatment model is correctly specified. Under correct treatment model specification, the Hybrid Inverse Probability of Treatment Weighted estimator is more efficient and has a lower mean squared standard error than its counterpart but does not maintain that advantage under treatment model misspecification at higher sample sizes.

With non-standard error terms generated using the student's t-distribution with 5 degrees of freedom, all of the advantages of the hybrid estimators at low sample sizes are still present but at larger sample sizes the hybrid methods faults are magnified in the case of treatment model misspecification. This is the result of the t-distribution being very tail heavy and this creates incorrect weights when the estimated treatment probability is poorly constructed and lead to very unstable estimates. One solution to this would be to truncate weights if, after modeling treatment assignment probabilities, probabilities were found that were very high or very low.

The sandwich estimator of the variance appears to be unbiased but unstable when the treatment model is misspecified. In particular the presence of the variable X_8 in the treatment models causes a large amount of bias. Removing the variable removes the bias.

However, when, in the case of our second simulation, the intercept is close in magnitude to the effect size of the treatments, the estimator becomes biased.

2.6 THE STAR*D TRIAL

The Sequenced Treatment Alternatives to Relieve Depression, or STAR*D, Trial was a 7 year multistage trial designed to test different regimes of depression treatments. The trial contained 4,041 patients between the ages of 18 and 75 who had received a clinical diagnosis of nonpsychotic major depressive disorder under the DSM-IV checklist. A patient that is considered to respond to treatment will continue taking that treatment, while those that do not respond are re-randomized to a treatment within a treatment-grouping of their own choice. Although the trial was designed with four and a half stages, we will only be looking at the 2nd and 3rd stage treat-ments. This is because everyone received the same treatment in the first stage, an SSRI called citalopram, and if we included the 4th stage into our analysis our regime sample sizes would be too small to make any reasonable inferences. We will also be excluding patients that received cognitive therapy in the second stage. They have the possibility of being assigned to an extra half stage, which would make analysis difficult in ways similar to if we had included the 4th stage in our analysis. So our analysis will center on comparing regimes for follow up medications for patients with depression

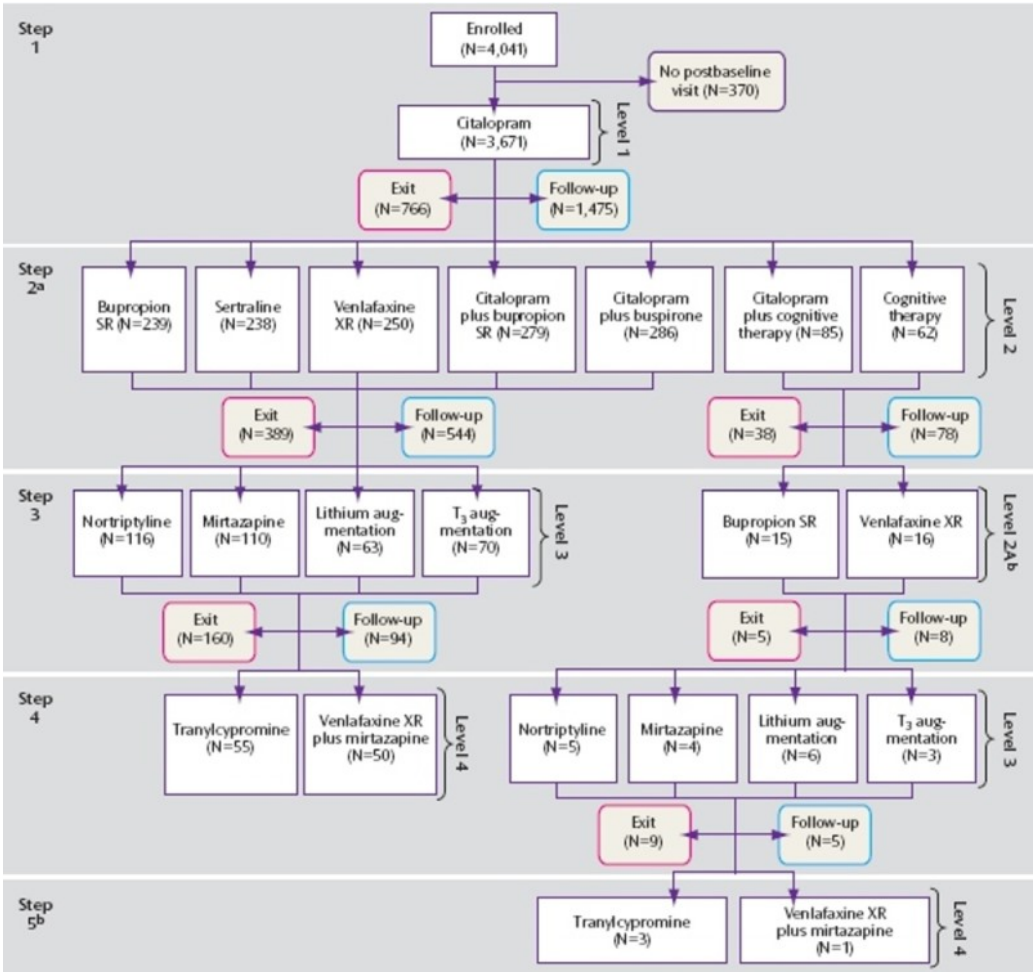

Figure 2: STAR*D Stages of Treatment.

Warden et al. [2007]

Table 11: STAR*D Dynamic Treatment Regimes

Regime	n	Mean (QIDS-16)	95% Confidence Interval
d(BUP,MIRT)	95	7.4	(7.3, 7.5)
d(BUP,NTP)	93	7.8	(7.7, 7.9)
d(BUP,BUP+LI)	85	7.7	(7.6, 7.9)
d(BUP,BUP+THY)	73	6.2	(5.9, 6.5)
d(SER,MIRT)	117	7.4	(7.4, 7.5)
d(SER,NTP)	108	6.7	(6.6, 6.8)
d(SER,SER+LI)	97	6.3	(6.2, 6.4)
d(SER,SER+THY)	96	6.1	(6.0, 6.2)
d(VEN,MIRT)	120	7.4	(7.3, 7.5)
d(VEN,NTP)	119	6.7	(6.6, 6.7)
d(VEN,VEN+LI)	108	6.9	(6.8, 7.0)
d(VEN,VEN+THY)	113	6.2	(6.2, 6.3)
d(CIT+BUP,MIRT)	157	6.4	(6.3, 6.5)
d(CIT+BUP,NTP)	163	6.7	(6.7, 6.8)
d(CIT+BUP,CIT+LI)	161	6.5	(6.4, 6.6)
d(CIT+BUP,CIT+THY)	167	5.9	(5.9, 6.0)
d(CIT+BUS,MIRT)	153	7.7	(7.7, 7.8)
d(CIT+BUS,NTP)	160	7.2	(7.1, 7.3)
d(CIT+BUS,CIT+LI)	145	8.1	(8.0, 8.2)
d(CIT+BUS,CIT+THY)	151	6.7	(6.7, 6.8)

depression that are non-responsive to SSRIs. For that reason stage 2 and 3 treatments will be henceforth be represented by 1st and 2nd stage notations under our conventions. So a response to stage 2 would be indicated by R_1, the outcome model for stage 3 would be denoted m_{jk}, and so on. A visual representation of the trial can be seen in figure 2.

Response is defined as either a 50% reduction in QIDS-16 (Quick Inventory of Depressive Symptomatology) self reported score, which the investigators call response, or a QIDS-16 score lower than 5, which the investigators call remission, at the end of the stage. The QIDS-16 score is also the outcome measure we will be looking at in our analysis, rather than the HRSD-17 (Hamilton Rating Scale of Depression) score which the original investigators used. This is because QIDS-16 was recorded at every visit while HRSD-17 was recorded upon exiting or finishing the trial, QIDS-16 scores would be available for more patients than the HRSD-17 was for. It should be noted that the QIDS-16 remission rates were generally higher than HRSD-17 rates, mainly because the study counted not having an exit HRSD-17 score as non-response. We will also differ from the study in that we will keep our overall outcome continuous rather than transform it to being binary. Transforming a continuous measure into a binary one is always accompanied by a loss of information. As with the original analysis by the investigators we assume, one could argue wrongly, that dropout from the study represents non-response and we use last observation carried forward as the final outcome.

The treatment probabilities for each stage were estimated using logistic regression modeling probability of receiving treatment of interest on the covariates Burden of Side Effects Rating (Wisniewski 2006) and change in QIDS-16 score that were measured during the previous stage. The outcomes were linearly modeled using our previously described β estimating technique with the covariates being age and QIDS-16 score at the beginning of stage 2.

Applying our new, more efficient doubly robust estimator to the data yields the results seen in the Tables and Figures below.

By this analysis, the best follow up treatment if citalopram is not effective is found to be either sertraline or citalopram combined with bupropion sustained release. Additionally the best regimes involve the third stage treatment of supplementing the second stage treatment with triiodothyronine. The best overall regime appears to be 'If non-response to citalopram augment it with bupropion sustained release and if non-response to that switch to supplementing citalopram with triiodothyronine.'

2.7 DISCUSSION

In this paper, we have introduced a new way of constructing estimators for two-stage dynamic treatment regimes, creating both a Hybrid Inverse Probability of Treatment Weighted estimator and a Hybrid Double Robust estimator. Through simulations we have shown that these estimators generally perform better than their non-hybrid counterparts by the metrics of bias, standard error, and mean standard error. Finally, we have applied the new Hybrid Double Robust estimator to the STAR*D data set in order to estimate the treatment effects of two stage dynamic treatment regimes for patients with non-psychotic major depressive disorder who do not respond to SSRIs.

One limitation of the methodology described in this chapter is that it discards data for whom the outcome data or treatment data are missing. This may result in bias, or efficiency loss as patients in trials can dropout for various reasons. Next chapter we introduce an estimator that, in addition to dealing with the missing data problem of not seeing the outcome under the regime of interest for all patients, addresses missing data due to dropout.

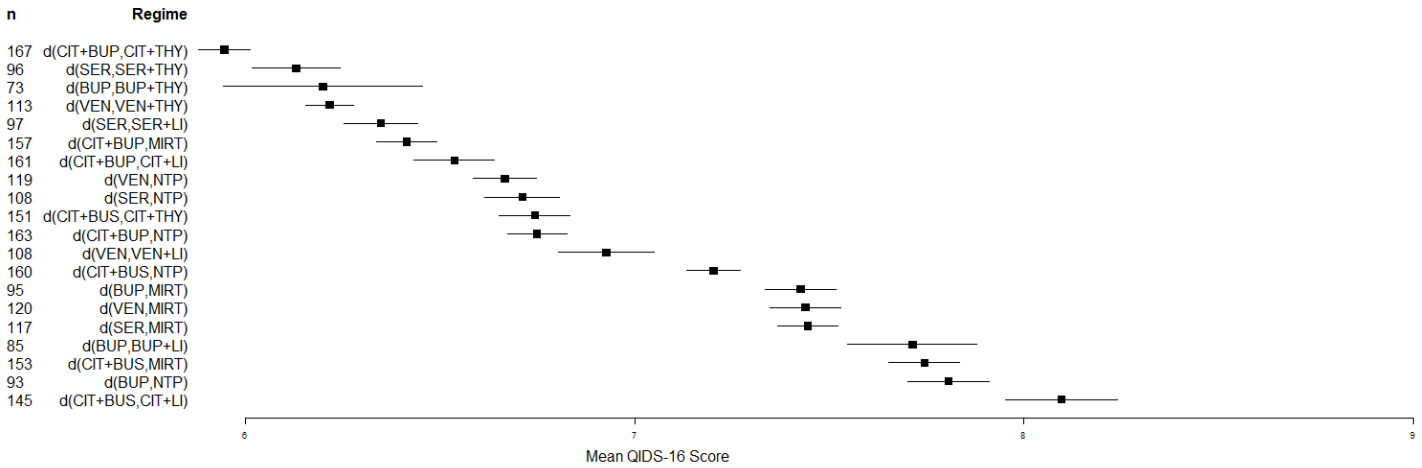

Figure 3: Forest Plot of STAR*D Dynamic Treatment Regimes

QIDS-16 Score of 5 = Remission; BUP - bupropion sustained release; BUS - buspirone; CIT - citalopram; Li - lithium; MIRT - mirtazapine; NTP - nortriptyline; SER - sertraline; THY - triiodothyronine

3.0 EFFICIENT INVERSE PROBABILITY OF TREATMENT WEIGHTING ESTIMATION FOR TWO-STAGE DYNAMIC TREATMENT REGIMES WITH DROP-OUT

3.1 INTRODUCTION

Drop-out is a common phenomenon in clinical studies. Patients will leave them for a variety of reasons and both legal and ethical constraints prevent investigators from collecting data from them after they dropout. Several issues that come up over the course of study that takes place over a long period may cause a lose of patients. Some reasons, such as a patient moving out of the area, are unrelated to the study itself and thus missing data stemming from them may be ignored. This type of missingness is typically called missing completely at random and analyses deleting missing data records remain unbiased. We call these complete case analyses. Other reasons, such as a patient leaving a study due to a lack of treatment efficacy or harsh side effects are most likely related to the treatments being administered within the study and thus ignoring this instance of missing data could result in an inaccurate analysis of the data. Statistically speaking, if the probability of a data point being observed given other observed data and missing data only depends on observed data, this type of missing data is called missing at random. This means the mechanism

causing dropout in such a scenario is observed. There are many methods designed to deal with data that are missing at random. When the probability of a data point being observed depends on the data that are unobserved, the mechanism is defined as missing not at random. Here we do not fully observe the mechanism of how the data are missing, such data is difficult to work with. This framework for missing data was first proposed by Rubin [1976]. Investigators spend a lot of time and effort on working with missing data in order to produce accurate and efficient results.

Studies comparing dynamic treatment regimes typically experience drop-out for a range of reasons. First, since the studies examine regimes of treatment, they can extend well past the length of a typical clinical trial. In such circumstances it is not unreasonable to assume that patients will experience a assortment of life events, such as moving, child birth, or even death, that could cause them to leave the study unexpectedly and for reasons mostly unrelated to the study itself. Second, if the dynamic treatment regimes are designed to find an efficacious treatment for a chronic disease, such as depression, then patients may become discouraged with the receipt of ineffective treatment and leave the study. Likewise dynamic treatment regimes for chronic illnesses such as cancer can include treatments with severe negative side effects that can cause a patient to leave the study. Thus it is important not to simply discard data on patients who fall into the second category, doing so can lead to biased and inefficient estimation of population parameters. From here on we will be assuming that all missing data are missing at random.

A common approach to missing data is to use observed data to impute the missing observations. Unconditional mean imputation, described by Little and Ruben [2002], uses the mean value of a variable across the observed data in place of any missing observation's value for that variable. Conditional mean imputation, defined by Buck[1960] also known as regression imputation or Buck's method, models a vari-

able, not seen for every observation, on a set of variables observed for all subjects. Then it uses that model to predict the missing data values. Hot deck imputation, originally developed for the Income Supplement of the Current Population Survey utilizes matching based on observed covariates in order to replace an observation with missing data values with values from observations without missing data.

Alternatively, one can use methods such as inverse probability weighting of complete cases by probability of being observed in order to account of missing data. The procedure is defined in detail in Chapters 1 & 2. A natural approach to dealing with missing data under a dynamic treatment regime setting is to use this approach, called inverse probability of treatment weighting (IPTW). Since we had previously discussed using IPTW to solve the missing data problem of not being on the regime of interest it is a natural course to use the same for drop out. The primary advantage of IPTW in a dynamic treatment regime setting as shown described by Bembom and van der Laan [2006] is that it allows for unbiased estimation of the outcome under a regime without having to know or specify the distribution of either the intermediate response or outcome. Defining an Inverse Probability of Treatment and Missing estimator is the first step toward presenting doubly robust or efficient estimators in a manner that takes missing data due to drop-out into account.

In this chapter we discuss causal inference in a dynamic treatment regime setting with data missing at random and more specifically, we describe two two IPTW estimators using different methods for estimating second stage treatment probability and derive the variance of these estimators. A simulation study then follows to compare all the discussed estimators in a multiplicity of conditions with other competing estimators. We then present a data analysis of the STAR*D clinical study in which we apply our estimator of interest to estimate the effects of two stage dynamic treatment regimes in the presence of missing data. Lastly, we will summarize the results in this chapter in a discussion section.

3.2 NOTATION

Consider the two-stage randomized trial considered in Chapter 2 where participants are assigned to a first stage treatment, with responders to that treatment being moved to follow-up and non-responders moved to a second stage of treatment. We will use similar notation as we did in Chapter 2, however we describe that again in detail here. For $i = 1, ..., n$ let $Z_{ji}^{(A)}$ be the indicator for the i^{th} patient receiving first stage treatment A_j, $Z_{ji}^{(A)} = 1$, if the patient receives A_j, 0, otherwise. Let R_i denote the indicator of whether participant i responds to the initial treatment. Conditional on not responding to the first stage treatment, let $Z_{ki}^{(B)}$ be the indicator for the i^{th} patient receiving second stage treatment B_k, $Z_{ki}^{(B)} = 1$, if the patient receives B_k, 0, otherwise. The observed data from the i^{th} participant for this trial in the presence

of missing data can be written as

$$O_i = [X^*_{1i}, \{Z^{(A)}_{ji}, j = 1, 2\}, \Delta^{(1)}_i, \Delta^{(1)}_i \{R_i, X^*_{2i}, (1 - R_i) Z^{(B)}_{ki} (k = 1, 2), \Delta^{(2)}_i, \Delta^{(2)}_i Y_i\}] \tag{3.1}$$

for $i = 1, ..., n$, where X^*_{1i} and X^*_{2i} are potential covariates prior to the randomization at stages 1 and 2, $\Delta^{(1)}_i = 1$ if the participant completed the first stage of treatment and stayed in the study until the second stage, 0 otherwise. $\Delta^{(2)}_i = 1$ if the participant's outcome Y_i was observed after responding to the first stage of treatment and being moved to follow-up or receiving a second stage of treatment after not responding to the first stage of treatment, 0 otherwise.

At each stage we observe covariate information, the i^{th} patient's observed covariate information for the first and second stages of treatment is X^*_{1i} and X^*_{2i}. Combined with the treatment assignments themselves we define $X_{1i} = (X^{*T}_{1i}, Z^{(A)}_{ji})^T$ and $X_{2i} = (X^{*T}_{2i}, R_i Z^{(B)}_{ki})^T$ as the observed data at stages 1 and 2. Define the i^{th} patient's outcome as Y_i.

For missing data define $\Delta^{(1)}_i$ and $\Delta^{(2)}_i$ as the i^{th} patient's indicators for not being missing at the first and second stages of treatment respectively. We define $\pi^{(1)}_i = Pr(\Delta^{(1)}_i = 1)$ and $\pi^{(2)}_i = Pr(\Delta^{(2)}_i = 1)$. It is seldom that these probabilities are inherently known in missing data problems. Based on the observed data at each stage we define the modeled probabilities of not being missing at each stage as $\pi^{(1)}(X_{1i}, \beta_1)$ and $\pi^{(2)}(X_{2i}, \beta_2)$, where β_1 and β_2 are the coefficients describing the relationship between the outcome and X_{1i} and X_{2i} respectively with a logit link.

Our goal is to estimate $\mu[d(A_j, B_k)] = E\{Y[d(A_j, B_k)]\}$, the mean of a population who were treated with $d(A_j, B_k)$. The regime 'Treat with A_j, if non-response treat with B_k'. For brevity's sake we will refer to $\mu[d(A_j, B_k)]$ from now on as simply μ.

3.3 CAUSAL INFERENCE FOR DYNAMIC TREATMENT REGIMES WITH DATA MISSING AT RANDOM

Consider a single stage scenario, the interest is to estimate the mean of an outcome Y, $\mu = E[Y|A_1]$ based on a sample of participants receiving A_1 and other treatments. The Inverse Probability of Treatment Weighted Estimator of μ is given by

$$\hat{\mu}_{IPTW} = \frac{1}{n}\sum_{i=1}^{n}\frac{Z_i^{(A)}}{\pi^{(A)}}Y_i \ , \tag{3.2}$$

where $Z_i^{(A)}$ is the indicator for receiving the treatment of interest in this one stage setting and $\pi^{(A)}$ is the corresponding probability of receiving that treatment.

This IPTW estimator solves the missing data problem of not observing the counterfactual outcome under the treatment of interest due to a portion of the sample receiving a treatment other than A_1. We can expand this missing data problem to include missing data due to dropout. Suppose that in addition to having patients whose outcome under A_1 is not observed due to receiving A_2, we also have patients drop out of the study after randomization and thus have missing outcome data. Consider the modified IPTW estimator

$$\hat{\mu}_{IPTMW} = \frac{1}{n}\sum_{i=1}^{n}\frac{\partial_i Z_i^{(A)}}{Pr(\Delta_i^{(1)} = 1|X_i, Z_i^{(A)})}Y_i \ , \tag{3.3}$$

where ∂_i indicates if the data are observed or not, $Pr(\Delta_i^{(1)} = 1|X_i, Z_i^{(A)})$ is the probability that the data is observed given the covariates X_i and $Z_i^{(A)}$ and $\hat{}$ indicates that this probability is estimated from the data. Now we have an estimator that utilizes weighting to solve two different missing data problems, that of not being on the treatment of interest as well as that of dropout.

When $Pr(\Delta_i^{(1)} = 1|X_i, Z_i^{(A)})$ is known we can see that $\hat{\mu}_{IPTMW}$ is an unbiased estimator of $\mu = E[Y|A_1]$

$$\begin{aligned}
E[\hat{\mu}_{IPTMW}] &= E\left[\frac{1}{n}\sum_{i=1}^{n}\frac{\partial_i Z_i^{(A)} Y_i}{\pi^{(A)} Pr(\Delta_i^{(1)} = 1|X_i, Z_i^{(A)})}\right] \\
&= E\left[\frac{\partial_i Z_i^{(A)} Y_i}{\pi^{(A)} Pr(\Delta_i^{(1)} = 1|X_i, Z_i^{(A)})}\right] \\
&= E\left[E\left[\frac{\partial_i Z_i^{(A)} Y_i}{\pi^{(A)} Pr(\Delta_i^{(1)} = 1|X_i, Z_i^{(A)})}\middle| Z_i^{(A)}, X_i\right]\right] \\
&= E\left[\frac{Z_i^{(A)}}{\pi^{(A)} Pr(\Delta_i^{(1)} = 1|X_i, Z_i^{(A)})} E\left[\partial_i Y_i \middle| Z_i^{(A)}, X_i\right]\right] \\
&= E\left[\frac{E[\partial_i|X_i, Z_i^{(A)}]\pi^{(A)}}{\pi^{(A)} Pr(\Delta_i^{(1)} = 1|X_i, Z_i^{(A)})} E\left[Y_i \middle| Z_i^{(A)} = 1, X_i\right]\right] \\
&= E[Y|A_j] \ .
\end{aligned} \tag{3.4}$$

Moving into a two stage dynamic treatment regime setting these exist for a variety of estimators for estimating mean outcome under a specific regime.

For complete data, Bembom and van der Laan [2006] describes a method, known as g-computation, originally described by Robins [1986], that uses weighted averages, based on response probability, of conditional intermediate responses to achieve unbiased estimates of mean regime outcomes. Specifically, the mean outcome under regime $d(A_j, B_k)$ is estimated as

$$\hat{\mu}_G[d(A_j, B_k)] = \frac{1}{n}\sum_{i=1}^{n}\left\{ R_i m_j(X_{1i}^*, X_{2i}^*) \ + (1 - R_i) m_{jk}(X_{1i}^*, X_{2i}^*) \right\}, \tag{3.5}$$

where $m_j(X_{1i}^*, X_{2i}^*)$ and $m_{jk}(X_{1i}^*, X_{2i}^*)$ respectively are the estimated mean outcome for responders receiving A_j on initial treatment and the same for non-responders

on treatment path $A_j \rightarrow B_k$. These can be estimated from the observed data. When the pathway means are correctly specified this estimator has very low bias and is extremely efficient. In the presence of missing data these quantities can be estimated using only complete cases but these may be biased estimates, which may cause the estimator itself to be biased.

When there are no missing data the Inverse Probability of Treatment Weighted estimator in a two stage dynamic treatment regime setting is described by Wahed and Tsiatis [2004] is

$$\hat{\mu}_{IPTW}[d(A_j, B_k)] = \frac{1}{n}\sum_{i=1}^{n} \left\{ \frac{Z_{ji}^{(A)}}{\pi_j^{(A)}} \left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\pi_{jk}^{(B)}} \right] \right\} Y_i \ . \tag{3.6}$$

Note that, following our previous arguments, under the assumption of consistency, this estimator is an unbiased estimator of $\mu[d(A_j, B_k)]$ when $\pi_j^{(A)}$ and $\pi_{jk}^{(B)}$ are known. We expand our IPTW estimator to include weights for missing via dropout. However, unlike the previous single stage estimator, there is now dropout during two different levels of treatment. This yields four different levels of missing data for the outcome under the regime of interest. First is for patients who receive a first level treatment other than that of interest, second is for patients who dropout during the first stage, third for patients who receive a second level treatment other than that of interest, and finally, for those that dropout during the second level of treatment or during follow-up after responding to the first level of treatment. This leads to the natural IPTW estimator

$$\hat{\mu}_{IPTMW}[d(A_j, B_k)] = \frac{1}{n}\sum_{i=1}^{n} \left\{ \frac{Z_{ji}^{(A)}\Delta_i^{(1)}\Delta_i^{(2)}}{\hat{\pi}_j^{(A)}\pi^{(1)}(X_{1i}, \hat{\beta}_1)\pi^{(2)}(X_{2i}, \hat{\beta}_2)} \left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\hat{\pi}_{jk}^{(B)}} \right] \right\} Y_i \ , \tag{3.7}$$

where the hats indicate that $\pi^{(1)}(X_{1i}, \hat{\beta}_1)$ and $\pi^{(2)}(X_{2i}, \hat{\beta}_2)$ are modeled estimators of $\pi_i^{(1)}$ and $\pi_i^{(2)}$ respectively based on the data prior to the missing. These models utilize logistic regression to express treatment probability as a function of the observed covariates. It is important to note that $\pi^{(2)}(X_{2i}, \hat{\beta}_2)$ is calculated from only those for who $\Delta_i^{(1)} = 1$. Likewise, $\hat{\pi}_{jk}^{(B)}$ and $\hat{\pi}_{jk}^{(B)}$ are estimates of $\pi_{jk}^{(B)}$ and $\pi_{jk}^{(B)}$ respectively, calculated as the mean treatment assignments at each level for the observed data.

This is our Inverse Probability of Treatment and Missing Weighted estimator and it will be the focus of this chapter.

3.4 VARIANCE OF THE IPTMW ESTIMATOR

3.4.1 When $\pi_{jk}^{(B)}$ is estimated using Complete Cases

The Inverse Probability of Treatment and Missing Weighted estimator described in the previous section is given by

$$\hat{\mu}_{IPTMW}[d(A_j, B_k)] = \frac{1}{n}\sum_{i=1}^{n}\left\{\frac{Z_{ji}^{(A)}\Delta_i^{(1)}\Delta_i^{(2)}}{\hat{\pi}_j^{(A)}\pi^{(1)}(X_{1i}, \hat{\beta}_1)\pi^{(2)}(X_{2i}, \hat{\beta}_2)}\left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\hat{\pi}_{jk}^{(B)}}\right]\right\}Y_i\ , \tag{3.8}$$

where $\hat{\beta}_1$ and $\hat{\beta}_2$ are estimated via logistic regression of $\Delta_i^{(1)}$ and $\Delta_i^{(2)}$ on the covariates observed prior to patients dropout. In other words $\hat{\beta}_1$ and $\hat{\beta}_2$ are solutions of the equations

$$0 = \sum_{i=1}^{n} X_{1i}\left(\Delta_i^{(1)} - \frac{1}{1+e^{-X_{1i}^T\beta_1}}\right) \tag{3.9}$$

and

$$0 = \sum_{i=1}^{n} \frac{\Delta_i^{(1)}}{\pi^{(1)}(X_{1i}, \beta_1)} X_{2i} \left(\Delta_i^{(2)} - \frac{1}{1 + e^{-X_{2i}^T \beta_2}} \right) . \tag{3.10}$$

In this situation, $\hat{\pi}_j^{(A)}$ and $\hat{\pi}_{jk}^{(B)}$ are estimated using the empirical average of participants receiving A_j and B_k after A_j respectively. In other words

$$\hat{\pi}_j^{(A)} = \frac{1}{n} \sum_{i=1}^{n} Z_{ji}^{(A)} \tag{3.11}$$

and

$$\hat{\pi}_{jk}^{(B)} = \frac{\sum_{i=1}^{n} \Delta_i^{(1)} \Delta_i^{(2)} Z_{ji}^{(A)} Z_{ki}^{(B)}}{\sum_{i=1}^{n} \Delta_i^{(1)} \Delta_i^{(2)} Z_{ji}^{(A)}} . \tag{3.12}$$

Putting these all together, we can express the IPTMW estimator as a solution of μ to the estimating equations

$$\sum_{i=1}^{n} \psi_i(O_i; \hat{\theta}) = \begin{pmatrix} \psi_{iA}(Z_{ji}^{(A)}; \hat{\pi}_j^{(A)}) \\ \psi_{i1}(X_{1i}, \Delta_i^{(1)}; \hat{\beta}_1) \\ \psi_{iB}(Z_{ki}^{(B)}, X_{1i}, \Delta_i^{(1)}; \hat{\pi}_{jk}^{(B)}) \\ \psi_{i2}(X_{2i}, \Delta_i^{(1)}, \Delta_i^{(2)}; \hat{\beta}_1, \hat{\beta}_2) \\ \psi_{i\mu}(O_i; \hat{\mu}) \end{pmatrix} = 0 \tag{3.13}$$

where

$$\theta = (\pi_j^{(A)}, \beta_1, \pi_{jk}^{(B)}, \beta_2, \mu) . \tag{3.14}$$

a$\hat{}$indicates a solution, and the individual ψ functions are the estimating equations corresponding to all the estimated parameters in the estimator $\hat{\mu}$ along with the estimate of μ itself. More specifically,

$$
\begin{aligned}
\psi_{iA}(Z_{ji}^{(A)};\pi_j^{(A)}) = & \quad Z_{ji}^{(A)} - \pi_j^{(A)} \\
\psi_{i1}(X_{1i},\Delta_i^{(1)};\beta_1) = & \quad X_{1i}\left(\Delta_i^{(1)} - \frac{1}{1+e^{-X_{1i}^T\beta_1}}\right) \\
\psi_{iB}(Z_{ki}^{(B)},X_{1i},\Delta_i^{(1)};\hat{\pi}_{jk}^{(B)}) = & \quad Z_{ji}^{(A)}(1-R_i)\Delta_i^{(1)}(Z_{ki}^{(B)} - \pi_{jk}^{(B)}) \\
\psi_{i2}(X_{2i},\Delta_i^{(1)},\Delta_i^{(2)};\beta_1,\beta_2) = & \quad \frac{\Delta_i^{(1)}}{\pi^{(1)}(X_{1i},\beta_1)} X_{2i}\left(\Delta_i^{(2)} - \frac{1}{1+e^{-X_{2i}^T\beta_2}}\right) \\
\psi_{i\mu}(O_i;\mu) = & \quad \left\{\frac{Z_{ji}^{(A)}\Delta_i^{(1)}\Delta_i^{(2)}}{\pi_j^{(A)}\pi^{(1)}(X_{1i},\beta_1)\pi^{(2)}(X_{2i},\beta_2)}\left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\pi_{jk}^{(B)}}\right]\right\}Y_i - \mu \ .
\end{aligned}
$$

By defining the above estimator as a solution to a set of estimating equations we are able to classify it as an M-Estimator. It follows then that it is asymptotically normal, as well as consistent (Stefanski and Boos [2002]) as long as the estimating equations are correctly specified. We can therefore use the sandwich estimator, V_n, to obtain an estimate of its asymptotic variance, namely

$$
V_n = A_n(\hat{\theta})^{-1}B_n(\hat{\theta})\{A_n(\hat{\theta})^{-1}\}^T/n \ , \tag{3.15}
$$

where

$$
A_n(\theta) = -\frac{1}{n}\sum_{i=1}^{n}\frac{\partial\psi_i(\theta)}{\partial\theta} \ , \tag{3.16}
$$

and

$$
B_n(\theta) = \frac{1}{n}\sum_{i=1}^{n}\psi_i(\theta)\psi_i(\theta)^T \ . \tag{3.17}
$$

More explicitly, the $A_n(\theta)$ can be written as

$$A_n(\theta) = \frac{1}{n}\sum_{i=1}^{n} -\frac{\partial \psi_i(\theta)}{\partial \theta} = \frac{1}{n}\sum_{i=1}^{n} \begin{pmatrix} -1 & 0 & 0 & 0 & 0 \\ 0 & -\frac{\partial \psi_{i1}}{\partial \beta_1} & 0 & 0 & 0 \\ 0 & 0 & -Z_{ji}^{(A)}(1-R_i)\Delta_i^{(1)} & 0 & 0 \\ 0 & -\frac{\partial \psi_{i2}}{\partial \beta_1} & 0 & -\frac{\partial \psi_{i2}}{\partial \beta_2} & 0 \\ -\frac{\partial \psi_{i\mu}}{\partial \pi_j^{(A)}} & -\frac{\partial \psi_{i\mu}}{\partial \beta_1} & -\frac{\partial \psi_{i\mu}}{\partial \pi_{jk}^{(B)}} & -\frac{\partial \psi_{i\mu}}{\partial \beta_2} & -1 \end{pmatrix}$$

where the derivative elements in the A matrix are as follows

$$\frac{\partial \psi_{i1}}{\partial \beta_1} = -\pi^{(1)}(X_{1i},\beta_1)(1-\pi^{(1)}(X_{1i},\beta_1))X_{1i}X_{1i}^T$$

$$\frac{\partial \psi_{i2}}{\partial \beta_1} = -(\Delta_i^{(2)} - \pi^{(2)}(X_{2i},\beta_2))X_{2i}X_{1i}^T\left(\frac{1}{\pi^{(1)}(X_{1i},\beta_1)} - 1\right)\Delta_i^{(1)}$$

$$\frac{\partial \psi_{i2}}{\partial \beta_2} = -\pi^{(2)}(X_{2i},\beta_2)(1-\pi^{(2)}(X_{2i},\beta_2))X_{2i}X_{2i}^T\frac{\Delta_i^{(1)}}{\pi^{(1)}(X_{1i},\beta_1)}$$

$$\frac{\partial \psi_{i\mu}}{\partial \pi_j^{(A)}} = -\left\{ \frac{Z_{ji}^{(A)}\Delta_i^{(1)}\Delta_i^{(2)}}{(\pi_j^{(A)})^2\pi^{(1)}(X_{1i},\beta_1)\pi^{(2)}(X_{2i},\beta_2)} \left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\pi_{jk}^{(B)}} \right] \right\} Y_i$$

$$\frac{\partial \psi_{i\mu}}{\partial \beta_1} = -\left\{ \frac{Z_{ji}^{(A)}\Delta_i^{(1)}\Delta_i^{(2)}}{\pi_j^{(A)}\pi^{(2)}(X_{2i},\beta_2)}\left(\frac{1}{\pi^{(1)}(X_{1i},\beta_1)} - 1\right) \left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\pi_{jk}^{(B)}} \right] \right\} X_{1i}^T Y_i$$

$$\frac{\partial \psi_{i\mu}}{\partial \pi_{jk}^{(B)}} = -\left\{ \frac{Z_{ji}^{(A)}(1-R_i)Z_{ki}^{(B)}\Delta_i^{(1)}\Delta_i^{(2)}}{\pi_j^{(A)}(\pi_{jk}^{(B)})^2\pi^{(1)}(X_{1i},\beta_1)\pi^{(2)}(X_{2i},\beta_2)} \right\} Y_i$$

$$\frac{\partial \psi_{i\mu}}{\partial \beta_2} = -\left\{ \frac{Z_{ji}^{(A)}\Delta_i^{(1)}\Delta_i^{(2)}X_{2i}}{\pi_j^{(A)}\pi^{(1)}(X_{1i},\beta_1)}\left(\frac{1}{\pi^{(2)}(X_{2i},\beta_2)} - 1\right) \left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\pi_{jk}^{(B)}} \right] \right\} X_{2i}^T Y_i$$

3.4.2 Using A Weighted Estimate Of $\pi_{jk}^{(B)}$

Previously we estimate $\pi_{jk}^{(B)}$ as the mean second stage treatment assignment using the observed data. However, this process of estimating $\pi_{jk}^{(B)}$ does not take into account patients who dropout out in the first stage. In order to recover information from patients who dropped out during the first stage, let us define a new estimator of $\pi_{jk}^{(B)}$,

$$\hat{\pi}_{jk}^{(B)} = \frac{\sum_{i=1}^{n} \frac{Z_{ji}^{(A)}(1-R_i)\partial_1 Z_{ki}^{(B)}}{\pi^{(1)}(X_{1i},\hat{\beta}_1)}}{\sum_{i=1}^{n} \frac{Z_{ji}^{(A)}(1-R_i)\partial_1}{\pi^{(1)}(X_{1i},\hat{\beta}_1)}} \ . \tag{3.18}$$

This estimator uses inverse probability weighting to account for patients whose second stage treatment assignment was never seen due to dropout. The Inverse Probability of Treatment and Missing Weighted estimator has not changed from (3.8) except that now $\pi_{jk}^{(B)}$ is estimated via (3.18) instead of (3.12). We will refer to this estimator as IPTMW2.

Similarly to in the previous section we can express the IPTMW with a weighted estimate of $\pi_{jk}^{(B)}$ as the solution to

$$\sum_{i=1}^{n} \psi_i(O_i;\hat{\theta}) = \begin{pmatrix} \psi_{iA}(Z_{ji}^{(A)};\hat{\pi}_j^{(A)}) \\ \psi_{i1}(X_{1i},\Delta_i^{(1)};\hat{\beta}_1) \\ \psi_{iB}(Z_{ki}^{(B)};\hat{\beta}_1,\hat{\pi}_{jk}^{(B)}) \\ \psi_{i2}(X_{2i},\Delta_i^{(1)},\Delta_i^{(2)};\hat{\beta}_1,\hat{\beta}_2) \\ \psi_{i\mu}(O_i;\hat{\mu}) \end{pmatrix} = 0 \tag{3.19}$$

where

$$O_i = (X_{1i}^*, X_{2i}^*, \Delta_i^{(1)}, \Delta_i^{(2)}, Z_{ji}^{(A)}, \Delta_i^{(1)}(1-R_i)Z_{ki}^{(B)}, \Delta_i^{(1)}R_i, Y_i) \text{ and } \theta = (\pi_j^{(A)}, \beta_1, \beta_2, \pi_B, \mu) \ . \tag{3.20}$$

Once again these equations are the estimating equations corresponding to all the estimated parameters in the estimator along with the estimate itself. The equations ψ_{iA}, ψ_{i1}, ψ_{i2}, $\psi_{i\mu}$ are as described before, and

$$\psi_{iB}(Z_{ki}^{(B)};\beta_1,\pi_{jk}^{(B)}) = \frac{Z_{ji}^{(A)}(1-R_i)\Delta_i^{(1)}}{\pi^{(1)}(X_{1i},\beta_1)}(Z_{ki}^{(B)}-\pi_{jk}^{(B)})$$

By defining the above estimator as a solution to a set of estimating equations we are once again able to classify it as an M-Estimator. It follows then that it is asymptotically normal, as well as consistent (Stefanski and Boos [2002]) as long as the estimating equations are correctly specified. We can therefore use the sandwich estimator, V_n, to obtain an estimate of its variance.

$$V_n = A_n(\hat{\theta})^{-1}B_n(\hat{\theta})\{A_n(\hat{\theta})^{-1}\}^T/n \ , \tag{3.21}$$

where

$$A_n(\theta) = -\frac{1}{n}\sum_{i=1}^{n}\frac{\partial\psi_i(\theta)}{\partial\theta} \ , \tag{3.22}$$

and

$$B_n(\theta) = \frac{1}{n}\sum_{i=1}^{n}\psi_i(\theta)\psi_i(\theta)^T \ . \tag{3.23}$$

Now the $A_n(\theta)$ becomes

$$A_n(\theta) = \frac{1}{n}\sum_{i=1}^{n} -\frac{\partial \psi_i(\theta)}{\partial \theta} = \frac{1}{n}\sum_{i=1}^{n} \begin{pmatrix} -1 & 0 & 0 & 0 & 0 \\ 0 & -\frac{\partial \psi_{i1}}{\partial \beta_1} & 0 & 0 & 0 \\ 0 & -\frac{\partial \psi_{iB}}{\partial \beta_1} & \frac{-Z_{ji}^{(A)}(1-R_i)\Delta_i^{(1)}}{\pi_i^{(1)}} & 0 & 0 \\ 0 & -\frac{\partial \psi_{i2}}{\partial \beta_1} & 0 & -\frac{\partial \psi_{i2}}{\partial \beta_2} & 0 \\ -\frac{\partial \psi_{i\mu}}{\partial \pi_j^{(A)}} & -\frac{\partial \psi_{i\mu}}{\partial \beta_1} & -\frac{\partial \psi_{i\mu}}{\partial \pi_{jk}^{(B)}} & -\frac{\partial \psi_{i\mu}}{\partial \beta_2} & -1 \end{pmatrix}$$

where the derivative elements in the A matrix are as follows

$$\frac{\partial \psi_{i1}}{\partial \beta_1} = -\pi^{(1)}(X_{1i},\beta_1)(1-\pi^{(1)}(X_{1i},\beta_1))X_{1i}X_{1i}^T$$

$$\frac{\partial \psi_{iB}}{\partial \beta_1} = -Z_{ji}^{(A)}(1-R_i)\Delta_i^{(1)}X_{1i}^T\left(\frac{1}{\pi_i^{(1)}}-1\right)(Z_{ki}^{(B)}-\pi_{jk}^{(B)})$$

$$\frac{\partial \psi_{i2}}{\partial \beta_1} = -(\Delta_i^{(2)}-\pi^{(2)}(X_{2i},\beta_2))X_{2i}X_{1i}^T\left(\frac{1}{\pi^{(1)}(X_{1i},\beta_1)}-1\right)\Delta_i^{(1)}$$

$$\frac{\partial \psi_{i2}}{\partial \beta_2} = -\pi^{(2)}(X_{2i},\beta_2)(1-\pi^{(2)}(X_{2i},\beta_2))X_{2i}X_{2i}^T\frac{\Delta_i^{(1)}}{\pi^{(1)}(X_{1i},\beta_1)}$$

$$\frac{\partial \psi_{i\mu}}{\partial \pi_j^{(A)}} = -\left\{ \frac{Z_{ji}^{(A)}\Delta_i^{(1)}\Delta_i^{(2)}}{(\pi_j^{(A)})^2\pi^{(1)}(X_{1i},\beta_1)\pi^{(2)}(X_{2i},\beta_2)} \left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\pi_{jk}^{(B)}} \right] \right\} Y_i$$

$$\frac{\partial \psi_{i\mu}}{\partial \beta_1} = -\left\{ \frac{Z_{ji}^{(A)}\Delta_i^{(1)}\Delta_i^{(2)}}{\pi_j^{(A)}\pi^{(2)}(X_{2i},\beta_2)}\left(\frac{1}{\pi^{(1)}(X_{1i},\beta_1)}-1\right) \left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\pi_{jk}^{(B)}} \right] \right\} X_{1i}^T Y_i$$

$$\frac{\partial \psi_{i\mu}}{\partial \pi_{jk}^{(B)}} = -\left\{ \frac{Z_{ji}^{(A)}(1-R_i)Z_{ki}^{(B)}\Delta_i^{(1)}\Delta_i^{(2)}}{\pi_j^{(A)}(\pi_{jk}^{(B)})^2\pi^{(1)}(X_{1i},\beta_1)\pi^{(2)}(X_{2i},\beta_2)} \right\} Y_i$$

$$\frac{\partial \psi_{i\mu}}{\partial \beta_2} = -\left\{ \frac{Z_{ji}^{(A)}\Delta_i^{(1)}\Delta_i^{(2)}X_{2i}}{\pi_j^{(A)}\pi^{(1)}(X_{1i},\beta_1)}\left(\frac{1}{\pi^{(2)}(X_{2i},\beta_2)}-1\right) \left[R_i + \frac{(1-R_i)Z_{ki}^{(B)}}{\pi_{jk}^{(B)}} \right] \right\} X_{2i}^T Y_i$$

A notable difference between this variance estimator and the variance estimator for IPTMW using the complete case estimator of $\pi_{jk}^{(B)}$ is $\frac{\partial \psi_B}{\partial \beta_1}$, which is now non-zero.

3.5 SIMULATION

We based our simulation on Kang and Schafer (2007) with a few notable exceptions. We assume a two stage design, with non-responders from the first stage being re-randomized in the second stage. In this simulation we concerned ourselves with the estimation of the mean counterfactual outcome $\mu\{Y[d(A_1, B_1)]\} = E\{Y[d(A_1, B_1)]\}$, the mean outcome for the patient population had everyone been on the regime $d(A_1, B_1)$. The indicators for A_1 and B_2, $Z_1^{(A)}$ and $Z_2^{(B)}$ were generated from a Bernoulli distribution with $\pi_1^{(A)} = \pi_{11}^{(B)} = 0.5$. Let $Z_2^{(A)} = 1 - Z_1^{(A)}$ and $Z_2^{(B)} = 1 - Z_1^{(B)}$. Intermediate response, R_i was generated as

$$R \sim BERNOULLI(\pi_R), \tag{3.24}$$

where $\pi_R = \tau_0 + \tau_1 * Z_1^{(A)}, j = 1, 2$. We consider various combinations of τ_0 and τ_1 to produce varying response rates. For example, $(\tau_0, \tau_1) = (0.2, 0.25)$ produces a 45% response rate among patients who receive treatment A_1.

The outcome variable, Y, is then generated as

$$Y = \gamma_0 + \gamma_1 * Z_1^{(A)} + \gamma_2 * Z_1^{(B)} * R + \gamma_3 * R + \gamma_4 * X_a + \gamma_5 X_b + e \ , \tag{3.25}$$

where X_a, X_b, and e are independent random variables generated from a standard normal distribution $[N(0, 1)]$ and $\gamma = (\gamma_0, \gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5)$ are varied to produce different relationships between the treatments, covariates, and outcomes. For $\gamma =$

$(7, 3, 2, 1, 5.5, 5.4)$, the true value of the estimand, the mean counterfactual outcome $\mu\{Y[d(A_1, B_1)]\} = E\{Y[d(A_1, B_1)]\}$ was calculated to be was 11.35.

The probability of being observed in the first stage was specified as

$$\pi^{(1)}(X_{1i}, \beta_1) = expit(\beta_{10} + \beta_{11} * X_a + \beta_{12} * Z_1^{(A)}) \ , \tag{3.26}$$

and $\Delta_i^{(1)}$ is then drawn from a Bernoulli distribution with probability $\pi^{(1)}(X_{1i}, \beta_1)$. Parameters $(\beta_{10}, \beta_{11}, \beta_{12})$ were varied to change the properties of missing data in the simulation; β_{11} is related to the strength of the association between missing data generation, probability of treatment assignment, and outcome, β_{12} controls the relative rate of missing data in the two groups, while β_{10} controls the amount of missing data which was varied to achieve between 10% and 50% of data being missing.

In the second stage the probability of being observed is defined as

$$\pi^{(2)} = expit(\beta_{20} + \beta_{21} * X_b + \beta_{22} * Z_j^{(A)} + \beta_{23} * Z_k^{(B)}) \tag{3.27}$$

where β_{20} was also varied to achieve a range of missing data between 10% and 50%. The values of β_{10}, β_{11}, and β_{20} along with the percentage of observations missing outcome data can be seen in Table 12. We performed a series of 5000 Monte Carlo simulations in order to asses the validity of our estimator under various levels of dropout and association between missing data generation, probability of treatment assignment, and outcome. We also varied sample size, running the simulation at $n = 500$ and $n = 2000$.

Table 12: Levels of Missing Data for varying β_{11} and β_{10} and β_{20}

β_{11}	β_{10}	β_{20}	Missing (%)
-.75	0.5	0.5	12%
-.75	1.5	1.5	26%
-.75	2.5	2.5	51%
-2.75	1	0.5	19%
-2.75	2	1.5	32%
-2.75	3	2.5	52%

$\beta_{11} = -.75$ corresponds to a strong association between probability of being missing, treatment assignment, and outcome and $\beta_{11} = -2.75$ corresponds to a weak one. $(\beta_{12}, \beta_{21}, \beta_{22}, \beta_{23})$ were fixed at (1.2, 0.4, -0.2, 1.3).

3.5.1 Simulation Results

In Table 13 we see the results of the simulation at $n = 500$ for 9% of observations missing for patients on the regime of interest and a strong association between probability of being missing, treatment assignment, and outcome. The Inverse Probability of Missing Weighted estimators have smaller relative bias (0.04% with complete case estimates of second stage treatment probability, IPTMW1, and 0.16% with the weighted estimate, IPTMW2) than both the complete case CC-G-Comp estimator, CC-G-Comp, (8.58%) and the complete case Inverse Probability of Treatment Weighted estimator, CC-IPTW, (15.1%). The Monte Carlo standard errors were smallest for the CC-G-Comp estimator (0.61) followed by IPTMW1 (0.77), CC-IPTW (0.81) and IPTMW2 (0.86). However, the composite metric of the MSE is smaller for the IPTMW estimators (1.20 and 1.48) than for both CC-IPTW (4.26) and the CC-G-Comp estimator (1.69). The standard errors obtained with bootstrapping largely echoed the Monte Carlo standard errors, more closely for CC-G-Comp (0.61 vs 0.61) and CC-IPTW (.81 vs .80) than for the IPTMW estimators (0.77 and 0.73 vs 0.86 and 0.79). Coverage probabilities for 95% confidence intervals constructed using the bootstrap standard errors yield better coverage for the IPTMW estimators (0.98% and 0.97%) than for CC-G-Comp (0.83%) and CC-IPTW (0.52%). This is more likely due to the large bias for the CC-G-Comp and CC-IPTW estimates.

Table 13 also shows that as the percentage of missing data increases (when there is a strong association between probability of being missing, treatment assignment, and outcome) that the IPTMW estimators continue to perform well across all metrics while CC-G-Comp and CC-IPTW tend to be more biased and inefficient. This is most obvious in their relative biases, which yield poor results at 32% of data missing

(13.9% and 25.5%) and abysmal results at 52% of data missing (21.5% and 41.8%). The coverage of 95% confidence intervals for CC-G-Comp and CC-IPTW at even moderate levels of missing data tell a similar story (50.0% and 13.7% coverage at 32% missing data).

The Table also shows the results at n=500 when there is only a weak association between missing data generation, probability of treatment assignment, and outcome. A notable finding is that CC-G-Comp yields good results at low and moderate levels of missing. At 26% missing data CC-G-Comp is only visibly worse then the IPTMW estimators in terms of relative bias (4.65% vs -0.10% and -0.06%), is almost similar in terms of coverage (94.4% vs 97.9% and 97.8%) and MSE (1.13 vs 1.04 and 1.12), and has a superior MCSE (0.65 vs 0.72 and 0.75).

Another interesting phenomenon is that across levels of missing data IPTMW1 has a higher relative bias than IPTMW2: at 12% missing data -0.07% vs -0.06%, at 26% missing data -0.10% vs -0.06%, and at 51% missing data -0.07% vs -0.04%.

All of these results are mirrored in Table 13 where the sample size is increased to n=2000 except that the variances and MSEs are lower because of the larger sample sizes.

In Table 15 we see how the estimate of the standard error obtained from the sandwich estimator of the variance compares to the Monte Carlo standard error at n=500 for when there are different levels association between missing data generation, probability of treatment assignment, and outcome. From the results in that Table we see that the estimates of the standard errors are more accurate when there is a weak association. These results are investigated for the larger sample size of n=2000, which can be seen in Table 16.

Table 13: n=500, Simulation results for estimating $\mu\{Y[d(A,B)]\} = 11.356$

Strength of Association	(%) Drop-Out	Method	Bias (%)	MCSE	MSE	Bootstrap SE	Coverage (%)
Strong	19%	$CC-G-Comp$	8.58	0.61	1.69	0.61	83
		$CC-IPTW$	15.10	0.81	4.26	0.80	52
		$IPTMW1$	0.04	0.77	1.20	0.73	98
		$IPTMW2$	0.16	0.86	1.48	0.80	97
	32%	$G-Comp$	13.90	0.64	3.30	0.64	50
		$IPTW$	25.50	0.95	10.18	0.95	14
		$IPTMW$	0.05	0.85	1.45	0.79	97
		$IPTMW2$	0.28	1.07	2.28	0.90	96
	52%	$G-Comp$	21.50	0.72	6.98	0.71	14
		$IPTW$	41.80	1.28	25.84	1.28	01
		$IPTMW$	0.23	1.02	2.08	0.93	95
		$IPTMW2$	0.71	1.31	3.45	1.08	94
Weak	12%	$G-Comp$	1.85	0.62	0.80	0.62	98
		$IPTW$	6.65	0.76	1.73	0.76	90
		$IPTMW$	-0.07	0.69	0.96	0.69	98
		$IPTMW2$	-0.06	0.70	0.98	0.70	98
	26%	$G-Comp$	4.65	0.65	1.13	0.65	94
		$IPTW$	16.10	0.90	4.98	0.91	52
		$IPTMW$	-0.10	0.72	1.04	0.72	98
		$IPTMW2$	-0.06	0.75	1.12	0.75	98
	51%	$G-Comp$	10.20	0.73	2.41	0.73	76
		$IPTW$	35.40	1.23	19.11	1.26	06
		$IPTMW$	-0.07	0.80	1.27	0.80	97
		$IPTMW2$	0.04	0.85	1.46	0.86	97

$CC-G-Comp$ = Complete G-Computation, $CC-IPTW$ = Complete Case Inverse Probability of Treatment Weighted Estimator, $IPTMW1$ = Inverse Probability of Missing Weighted Estimator with complete case estimator of $\pi_{jk}^{(B)}$, $IPTMW2$ = Inverse Probability of Missing Weighted Estimator with weighted estimate of $\pi_{jk}^{(B)}$

Table 14: n=2000, Simulation results for estimating $\mu\{Y[d(A,B)]\} = 11.356$

Strength of Association	(%) Drop-Out	Method	Bias (%)	MCSE	MSE	Bootstrap SE	Coverage (%)
Strong	19%	$G-Comp$	8.69	0.31	1.16	0.305	63
		$IPTW$	15.10	0.40	3.27	0.40	11
		$IPTMW$	-0.002	0.37	0.28	0.37	99.9
		$IPTMW2$	0.06	0.47	0.44	0.41	99.7
	32%	$G-Comp$	14.00	0.32	2.73	0.32	6
		$IPTW$	25.50	0.47	8.83	0.47	<0.1
		$IPTMW$	0.05	0.42	0.35	0.40	99.9
		$IPTMW2$	0.18	0.55	0.60	0.48	99.4
	51%	$G-Comp$	21.50	0.36	6.20	0.36	<0.1
		$IPTW$	41.60	0.63	23.10	0.64	<0.1
		$IPTMW$	0.09	0.53	0.57	0.49	99.6
		$IPTMW2$	0.28	0.69	0.96	0.59	98.8
Weak	11%	$G-Comp$	1.97	0.31	0.24	0.31	99.7
		$IPTW$	6.68	0.38	0.87	0.38	88
		$IPTMW$	-0.03	0.35	0.24	0.35	100
		$IPTMW2$	-0.03	0.35	0.25	0.35	99.9
	26%	$G-Comp$	4.79	0.33	0.51	0.33	95
		$IPTW$	16.20	0.45	3.79	0.45	12
		$IPTMW$	-0.02	0.36	0.27	0.36	100
		$IPTMW2$	-0.01	0.38	0.28	0.37	99.9
	51%	$G-Comp$	10.20	0.37	1.63	0.37	52
		$IPTW$	35.30	0.62	16.79	0.62	<0.1
		$IPTMW$	-0.01	0.40	0.32	0.40	99.8
		$IPTMW2$	0.01	0.43	0.37	0.43	99.7

$CC-G-Comp$ = Complete G-Computation, $CC-IPTW$ = Complete Case Inverse Probability of Treatment Weighted Estimator, $IPTMW1$ = Inverse Probability of Missing Weighted Estimator with complete case estimator of $\pi_{jk}^{(B)}$, $IPTMW2$ = Inverse Probability of Missing Weighted Estimator with weighted estimate of $\pi_{jk}^{(B)}$

Table 15: n=500, $\mu\{Y[d(A,B)]\} = 11.356$

Association	Level of Missing Data Overall	Level of Missing Data On Regime	Method	Sandwich Est. SE	MCSE	Coverage (%)
Strong	19%	19%	$IPTMW1$	0.774	0.722	95.1
			$IPTMW2$	0.859	0.776	94.9
	32%	31%	$IPTMW1$	0.853	0.797	95.3
			$IPTMW2$	1.070	0.883	95.1
	52%	48%	$IPTMW1$	1.020	0.967	95.9
			$IPTMW2$	1.310	1.070	95.1
Weak	12%	8%	$IPTMW1$	0.691	0.691	95.3
			$IPTMW2$	0.701	0.701	95.2
	26%	19%	$IPTMW1$	0.722	0.725	94.9
			$IPTMW2$	0.747	0.749	95.1
	51%	40%	$IPTMW1$	0.797	0.817	95.1
			$IPTMW2$	0.854	0.869	95.2

$IPTMW1$ = Inverse Probability of Missing Weighted Estimator with complete case estimator of $\pi_{jk}^{(B)}$, $IPTMW2$ = Inverse Probability of Missing Weighted Estimator with weighted estimate of $\pi_{jk}^{(B)}$

Table 16: n=2000, $\mu\{Y[d(A,B)]\} = 11.356$

Association	Level of Missing Data Overall	Level of Missing Data On Regime	Method	Sandwich Est. SE	MCSE	Coverage (%)
Strong	19%	19%	$IPTMW1$	0.372	0.369	95.0
			$IPTMW2$	0.471	0.410	94.9
	32%	28%	$IPTMW1$	0.421	0.417	95.8
			$IPTMW2$	0.548	0.478	95.1
	52%	47%	$IPTMW1$	0.534	0.539	96.8
			$IPTMW2$	0.693	0.615	95.8
Weak	11%	9%	$IPTMW1$	0.348	0.346	95.0
			$IPTMW2$	0.353	0.351	95.1
	26%	20%	$IPTMW1$	0.364	0.362	95.0
			$IPTMW2$	0.377	0.374	95.0
	51%	41%	$IPTMW1$	0.403	0.409	95.2
			$IPTMW2$	0.430	0.434	95.2

$IPTMW1$ = Inverse Probability of Missing Weighted Estimator with complete case estimator of $\pi_{jk}^{(B)}$, $IPTMW2$ = Inverse Probability of Missing Weighted Estimator with weighted estimate of $\pi_{jk}^{(B)}$

3.6 CONCLUSION

For a strong association between probability of being missing, treatment assignment, and outcome the IPTMW estimators are less biased, have a better MSE, and better coverage probabilities then both CC-IPTW and CC-G-Comp. As the level of missing data increases the IPTW and G-Computation estimators become inefficient and more biased than the IPTMW estimators.

For a weak association between probability of being missing, treatment assignment, and outcome at both sample sizes CC-G-Comp is competitive with both IPTMW estimators and even has a noticeably lower MCSE and MSE with low levels of missing data.

Across all levels of missing data the estimator of the variance of the IPTMW estimators is more accurate when there is a weak association between probability of being missing, treatment assignment, and outcome as opposed to a strong one. This may be caused by the increased variability of dropout in the first stage that exists as a result of manner in which we strengthen the association between probability of being missing, treatment assignment, and outcome in our simulation.

3.7 ANALYSIS OF STAR*D DATA

The Sequenced Treatment Alternatives to Relieve Depression, or STAR*D, Trial was a multistage trial designed to test different regimes of depression treatments. The trial contained 4,041 patients between the ages of 18 and 75 who had received a clinical diagnosis of nonpsychotic major depressive disorder under the DSM-IV checklist. A patient that is considered to respond to treatment will continue taking

that treatment, while those that do not respond are re-randomized to a treatment within a treatment-grouping of their own choice.

The trial was designed with 4 (and a half) stages. This study contains a unique element called equipoise stratified randomization, meaning that patients were given the option of selecting classes of treatments that they would find acceptable to be on and then were randomized to one of the treatments within those groups. This was done in order to boost compliance and decrease drop-out. A patient with more control over their treatment is thought to be more inclined to follow that treatment and not leave the study. Everyone received the same treatment in the first stage, an SSRI called citalopram. Non-responders at each stage were offered a choice of switching off of their current medication or supplementing it with another treatment. Non-responders to the second stage treatment who received cognitive therapy or citalopram supplemented with cognitive therapy that stage were randomized to one of two medicinal treatments at what investigators called stage 2A. The purpose of this was to make sure everyone that made it to the third stage received a follow up medication to citalopram. In the STAR*D Trial responders to third stage treatment are kept on their treatment while non-responders receive a fourth stage of treatment. However, as previously stated our analysis stops after the third stage. A visual representation of the trial can be seen in figure 4

Response between stages was defined as either a 50% reduction in QIDS-16 (Quick Inventory of Depressive Symptomatology) self reported score, or a QIDS-16 score lower than 5, at the end of the stage. The overall outcome of interest for investigators was the HRSD-17 (Hamilton Rating Scale of Depression) score which was collected upon study exit.

Although the trial was designed with 4 and a half stages we will only be looking at the 2nd and 3rd stage treatments. This is because everyone received the

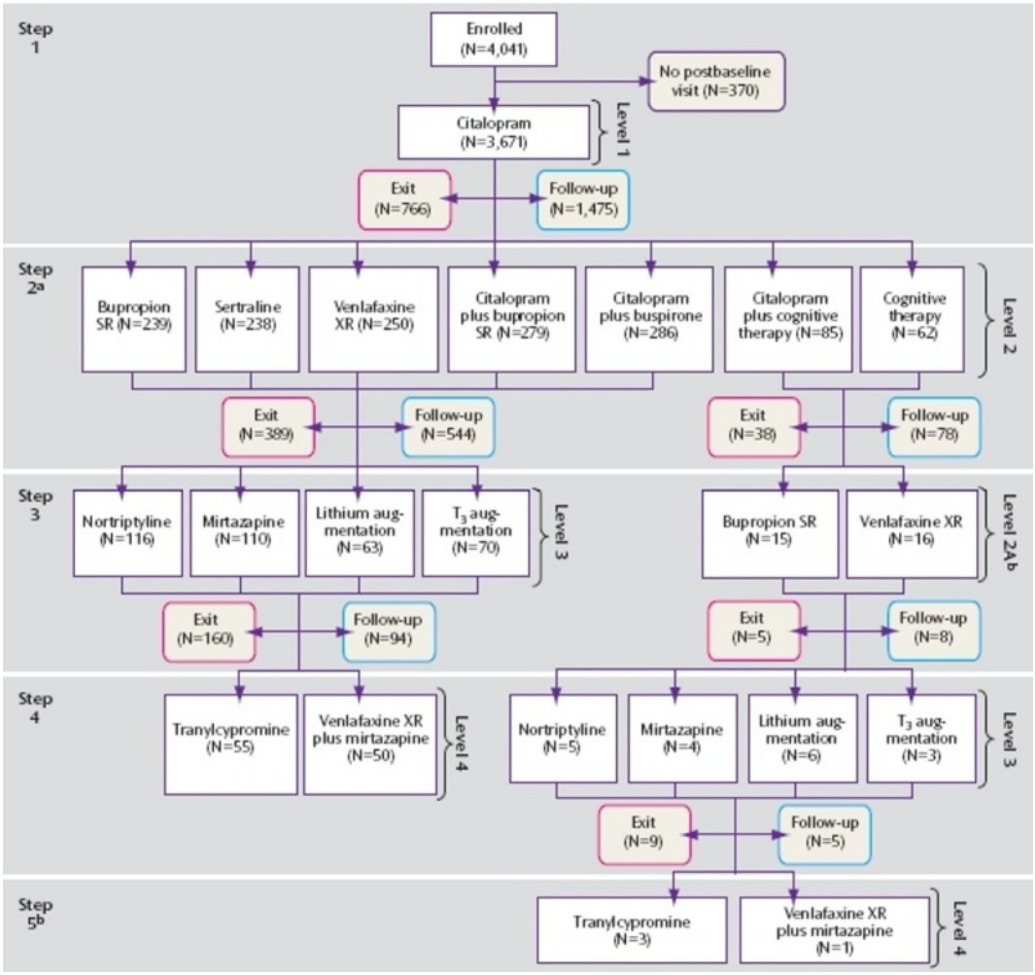

Figure 4: STAR*D Stages of Treatment.

Warden et al. [2007]

Table 17: STAR*D Dynamic Treatment Regimes

Regime	n	Mean (QIDS-16)	95% Confidence Interval
d(BUP,MIRT)	96	8.4	(8.0 , 8.7)
d(BUP,NTP)	97	9.1	(8.8 , 9.4)
d(BUP,BUP+LI)	86	9.71	(8.8 , 10.5)
d(BUP,BUP+THY)	76	4.6	(3.6 , 5.7)
d(SER,MIRT)	120	8.9	(8.6 , 9.2)
d(SER,NTP)	109	7.7	(7.4 , 8.0)
d(SER,SER+LI)	98	10.2	(9.5 , 10.9)
d(SER,SER+THY)	97	9.1	(8.2 , 9.9)
d(VEN,MIRT)	121	8.3	(8.0 , 8.6)
d(VEN,NTP)	122	6.4	(6.3 , 6.6)
d(VEN,VEN+LI)	110	7.1	(6.6 , 7.7)
d(VEN,VEN+THY)	115	7.1	(6.5 , 7.7)
d(CIT+BUP,MIRT)	160	5.5	(5.4, 5.6)
d(CIT+BUP,NTP)	167	6.2	(6.1, 6.4)
d(CIT+BUP,CIT+LI)	164	7.4	(7.0, 7.9)
d(CIT+BUP,CIT+THY)	170	6.9	(6.6 , 7.2)
d(CIT+BUS,MIRT)	154	5.8	(5.7, 5.9)
d(CIT+BUS,NTP)	162	6.8	(6.7 , 7.0)
d(CIT+BUS,CIT+LI)	146	8.7	(8.2 , 9.3)
d(CIT+BUS,CIT+THY)	153	7.3	(7.0 , 7.6)

same treatment in the first stage, an SSRI called citalopram, and if we included the 4th stage into our analysis our regime sample sizes would be too small to make any reasonable inferences. For the same reason we will also be excluding patients that received cognitive therapy in the second stage. So our analysis will center on comparing regimes for follow up medications for patients with depression that are non-responsive to SSRIs. For that reason stage 2 and 3 treatments will be henceforth be represented by 1st and 2nd stage notations under methodology developed in this chapter. So a response to stage 2 would be indicated by R, the indicator for being observed during stage 3 would be denoted Δ_2, and so on.

The QIDS-16 score is the outcome measure of interest in our analysis, rather than the HRSD-17 (Hamilton Rating Scale of Depression) score used as the primary outcome for the study. This is because QIDS-16 was recorded at every visit while HRSD-17 was recorded upon exiting or finishing the trial, QIDS-16 scores would be available for more patients than the HRSD-17 was for. It should be noted that the QIDS-16 remission rates were generally higher than HRSD-17 rates, mainly because the study counted not having an exit HRSD-17 score as non-response. We will also differ from the study in that we will keep our overall outcome continuous rather than transform it to being binary. Transforming a continuous measure into a binary one is always accompanied by a loss of information. For simplicity, we will be taking sample mean estimates, weighted on probability of being observed, of the probabilities of treatment.

The probabilities of being observed for each stage were estimated using logistic regression using variables indicated as being related to dropout by Warden et al. [2007]. The model for probability of not dropping out in the initial stage of treatment included age and treatment assigned at that stage as covariates. The model for probability of not dropping out in the second stage of treatment included a binary

variable for being non-white hispanic, treatment at the prior stage, and treatment assigned at that stage. Probability of receiving a particular treatment at each stage was calculated as it was in the simulation, using the mean of the indicators of treatment assignment at each stage.

Applying our new, IPTMW estimator with the complete case estimates of second stage treatment probabilities to the data yields the results seen in the tables and plots below. Applying the IPTMW2 estimator to the data yielded identical results. From Table 17 and Figure 5 we can infer that the best induction treatment is switch to bupropion sustained release and the best follow up treatment is augmenting with triiodothyronine. The wide confidence interval for $d(BUP, BUP{+}THY)$ is the result of a very small number of patients being on the treatment pathway $SER \rightarrow SER+THY$. Overall, the analysis indicated the best regime for patients not responding to citalopram is to augment with bupropion sustained release and if non-response to that augment instead with triiodothyronine this is contrary to in Chapter 2 that favors the same regime but augmenting with rather than switching to bupropion sustained release. Other differences between these results and those presented in Chapter 2 lay mostly in the size of the confidence intervals and in Chapter 2's estimates of regime means are mostly clustered around a QIDS-16 score of 7 while this Chapter has a larger range of values between scores of 4 and 11.

3.8 DISCUSSION

In this paper we have introduced and discussed estimators that account for participant drop-out in two stage dynamic treatment regimes. A new two stage Inverse Probability of Treatment and Missing Weighted estimator was defined, both utiliz-

ing a complete case estimate of second stage treatment assignment probability and an estimate based on weighting observed cases by the inverse of the probability of being observed during the previous stage. These estimators were compared to complete case estimators typically used to analyze two stage dynamic treatment regimes. Through these simulations we showed that in a variety of settings these new estimators were superior. Finally, we applied the Inverse Probability of Treatment and Missing Weighted estimator with a complete case estimate of second stage treatment assignment probability in order to estimate the treatment effects of different two stage dynamic treatment regimes for patients with non-psychotic major depressive disorder who do not respond to SSRIs in the presence of missing data.

In Chapter 3, we discussed the IPTMW estimator when treatment assignment probabilities were fixed. However, in observational studies, the probability patients receive a particular treatment is not fixed in advance and hence needs to be estimated based on patient's characteristics. Furthermore, following procedures similar to the one described in Chapter 2, we can augment the IPTMW estimator, to make it robust to misspecification of the dropout models.

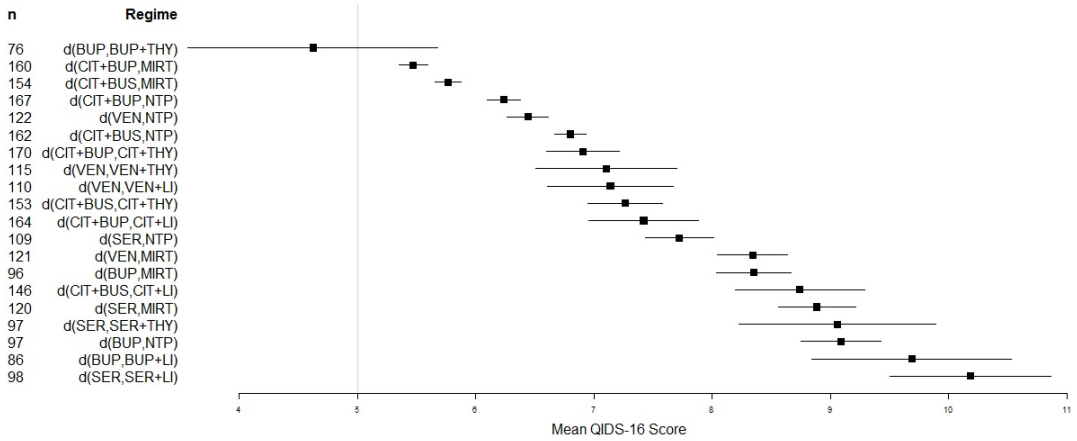

Figure 5: Forest Plot of STAR*D Dynamic Treatment Regimes

QIDS-16 Score of 5 = Remission; BUP - bupropion sustained release; BUS - buspirone; CIT - citalopram; Li - lithium; MIRT - mirtazapine; NTP - nortriptyline; SER - sertraline; THY - triiodothyronine

4.0 AN AUGMENTED INVERSE PROBABILITY OF TREATMENT AND MISSING WEIGHTED ESTIMATOR FOR TWO-STAGE DYNAMIC TREATMENT REGIMES

4.1 INTRODUCTION

Observational studies naturally generate data for comparing dynamic treatment regimes. The sample size, time required, and variety of treatments sometimes make conducting randomized clinical trials challenging. Then, even when clinical trials are conducted to study different dynamic treatment regimes they make take on aspects of observational studies such as non random treatment assignment. One example is the Sequenced Treatment Alternatives to Relieve Depression, or STAR*D, Trial, a multistage trial designed to test different regimes of depression treatments. In order to increase patient retention within the trial the investigators employed a technique called equipoise stratified randomization, where participants would select different groups of treatment that they found acceptable and then were randomized within those groups. This was done to give participants more control over the treatments they received, which in turn was thought to give them more incentive to adhere to those treatments. Thus the probability of being assigned to a given treatment at each level was neither known nor random.

Consider the same two-stage study set-up considered in Chapter 2 where participants receive a first stage treatment, with responders to that treatment being moved to follow-up and non-responders receiving a second stage of treatment except that the treatments are not randomly assigned. For $i = 1, ..., n$ let $Z_{ji}^{(A)}$ be the indicator for the i^{th} patient receiving first stage treatment A_j, $Z_{ji}^{(A)} = 1$, if the patient receives A_j, 0, otherwise. Let R_i denote the indicator of whether participant i responds to the initial treatment. Conditional on not responding to the first stage treatment, let $Z_{ki}^{(B)}$ be the indicator for the i^{th} patient receiving second stage treatment B_k, $Z_{ki}^{(B)} = 1$, if the patient receives B_k, 0, otherwise. The observed data can be written as

$$O_i = [X_{1i}^*, \{Z_{ji}^{(A)}, j = 1, 2\}, \Delta_i^{(1)}, \Delta_i^{(1)}\{R_i, X_{2i}^*, (1 - R_i)Z_{ki}^{(B)}(k = 1, 2), \Delta_i^{(2)}, \Delta_i^{(2)}Y_i\}] \tag{4.1}$$

for $i = 1, ..., n$, where X_{1i}^* and X_{2i}^* are potential covariates prior to the randomization at stages 1 and 2, $\Delta_i^{(1)} = 1$ if the participant completed the first stage of treatment and stayed in the study until the second stage, 0 otherwise. $\Delta_i^{(2)} = 1$ if the participant's outcome Y_i was observed after responding to the first stage of treatment and being moved to follow-up or receiving a second stage of treatment after not responding to the first stage of treatment, 0 otherwise.

At each stage we observe covariate information, the i^{th} patient's observed covariate information for first and second stages of treatment which we define as X_{1i}^* and X_{2i}^*. Combined with the treatment assignments themselves we define $X_{1i} = (X_{1i}^{*T}, Z_{ji}^{(A)})^T$ and $X_{2i} = (X_{2i}^{*T}, R_i Z_{ki}^{(B)})^T$ as the observed data prior to treatment assignment at stages 1 and 2.

We define $\pi_i^{(1)} = Pr(\Delta_i^{(1)} = 1)$ and $\pi_i^{(2)} = Pr(\Delta_i^{(2)} = 1)$. It is seldom that these probabilities are inherently known and we rely on models to estimate them. The modeled probabilities of not being missing at each stage are $\pi^{(1)}(X_{1i}, \beta_1)$ and $\pi^{(2)}(X_{2i}, \beta_2)$, where β_1 and β_2 are the coefficients describing the relationship between the outcome and X_{1i} and X_{2i} respectively.

Here the treatment assignment probabilities are not known, let us define them as $\pi_{ji}^{(A)} = Pr(Z_{ji}^{(A)} = 1)$ and $\pi_{jki}^{(B)} = Pr(Z_{ki}^{(B)} = 1 | R_i = 1)$. They are estimated using logistic regression models that include patient characteristics collected prior to treatment assignment at each stage.

Our goal is to estimate $\mu[d(A_j, B_k)] = E\{Y[d(A_j, B_k)]\}$, the mean of a population who were treated with $d(A_j, B_k)$. The regime 'Treat with A_j, if non-response treat with B_k. For brevity's sake we will refer to $\mu[d(A_j, B_k)]$ from now on as simply μ.'

In Chapter 3 we defined a two stage Inverse Probability of Treatment and Missing Weighted (IPTMW) estimator. However, that estimator assumed that probability treatment assignment at each stage was fixed given prior treatment assignment and response. So for a two-stage dynamic treatment regime in an observational setting that experiences missing data problems let us extend the IPTMW to

$$\hat{\mu}_{IPTMW}[d(A_j, B_k)] = \frac{1}{n} \sum_{i=1}^{n} \left\{ \frac{Z_{ji}^{(A)} \Delta_i^{(1)} \Delta_i^{(2)}}{\pi_j^{(A)}(X_{1i}, \hat{\gamma_1}) \pi^{(1)}(X_{1i}, \hat{\beta}_1) \pi^{(2)}(X_{2i}, \hat{\beta}_2)} \left[R_i + \frac{(1 - R_i) Z_{ki}^{(B)}}{\pi_{jk}^{(B)}(X_{2i}, \hat{\gamma_2})} \right] \right\} Y_i \ , \tag{4.2}$$

where β_1 and β_2 are parameters associated with the covariates that link them via the logit function to the indicator for being observed in the first and second stages of treatment respectively and γ_1 and γ_2 are perform the same function for the indicators for first and second stage treatment assignment.

Defining this Observational Inverse Probability of Treatment and Missing Weighted estimator is the first step toward moving the estimator introduced in Chapter 3 into an observational setting. This estimator makes use of patient characteristics to estimate probability of treatment assignment at each stage as opposed to the IPTMW estimator from Chapter 3 that assumed the probability was fixed across the study population conditional on prior treatment and response. The next step would be confirming this estimator is unbiased when the true model parameters are used and consistent when they are estimated. After that estimation of the variance would be necessary, likely once again using the sandwich estimator. A simulation would then check if this estimator performs better than the G-Computation and Inverse Probability of Treatment Weighted Estimators for observational data in the presence data missing due to drop-out wile also confirming that the estimator of the variance is unbiased. Finally this new estimator would be applied to the STAR*D study in order to utilize data for patients that drop-out while on the regimes of interest while accounting for the equipoise stratified randomization present in the study.

BIBLIOGRAPHY

[1] Douglas G Altman and Patrick Royston. The cost of dichotomising continuous variables. *Bmj*, 332(7549):1080, 2006.

[2] Xiaofei Bai, Anastasios A Tsiatis, and Sean M O'Brien. Doubly-robust estimators of treatment-specific survival distributions in observational studies with stratified sampling. *Biometrics*, 69(4):830–839, 2013.

[3] Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.

[4] Oliver Bembom and Mark J van der Laan. Statistical methods for analyzing sequentially randomized trials. *Journal of the National Cancer Institute*, 99(21):1577–1582, 2007.

[5] Jolene Birmingham, Andrea Rotnitzky, and Garrett M Fitzmaurice. Pattern–mixture and selection models for analysing longitudinal data with monotone missing patterns. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(1):275–297, 2003.

[6] Samuel F Buck. A method of estimation of missing values in multivariate data suitable for use with an electronic computer. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 302–306, 1960.

[7] Weihua Cao, Anastasios A Tsiatis, and Marie Davidian. Improving efficiency and robustness of the doubly robust estimator for a population mean with incomplete data. *Biometrika*, page asp033, 2009.

[8] James R Carpenter, Michael G Kenward, and Stijn Vansteelandt. A comparison of multiple imputation and doubly robust estimation for analyses with missing

data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 169(3):571–584, 2006.

[9] Miguel A Hernán and James M Robins. Estimating causal effects from epidemiological data. *Journal of epidemiology and community health*, 60(7):578–586, 2006.

[10] Paul W Holland. Statistics and causal inference. *Journal of the American statistical Association*, 81(396):945–960, 1986.

[11] Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952.

[12] Joseph DY Kang and Joseph L Schafer. Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical science*, pages 523–539, 2007.

[13] Jared K Lunceford, Marie Davidian, and Anastasios A Tsiatis. Estimation of survival distributions of treatment policies in two-stage randomization designs in clinical trials. *Biometrics*, 58(1):48–57, 2002.

[14] Katherine K Matthay, Judith G Villablanca, Robert C Seeger, Daniel O Stram, Richard E Harris, Norma K Ramsay, Patrick Swift, Hiroyuki Shimada, C Thomas Black, Garrett M Brodeur, et al. Treatment of high-risk neuroblastoma with intensive chemotherapy, radiotherapy, autologous bone marrow transplantation, and 13-cis-retinoic acid. *New England Journal of Medicine*, 341(16):1165–1173, 1999.

[15] Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355, 2003.

[16] Susan A Murphy. An experimental design for the development of adaptive treatment strategies. *Statistics in medicine*, 24(10):1455–1481, 2005.

[17] Mitsuo Ono and Herman P Miller. *Income nonresponses in the current population survey*. US Bureau of the Census, 1969.

[18] James M Robbins, Miguel A Hernan, and Babette Brumback. Marginal structural models and causal inferences in epidemiology. *Epidemiology*, 11:550–60, 2000.

[19] James Robins. A new approach to causal inference in mortality studies with a sustained exposure periodapplication to control of the healthy worker survivor effect. *Mathematical Modelling*, 7(9):1393–1512, 1986.

[20] James M Robins. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association*, volume 1999, pages 6–10, 2000.

[21] James M Robins and Andrea Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(429):122–129, 1995.

[22] James M Robins, Andrea Rotnitzky, and Daniel O Scharfstein. Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models. In *Statistical models in epidemiology, the environment, and clinical trials*, pages 1–94. Springer, 2000.

[23] James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association*, 89(427):846–866, 1994.

[24] JM Robins. Marginal structural models versus nested structural models as tools for causal inference. *Statistical models in epidemiology (Edited by: Halloran E) New York, Springer Verlag*, 1999.

[25] JM Robins and A Rotnitzky. Årecovery of information and adjustment for dependent censoring using surrogate markers. g in: Aids epidemiology m methodological issues. eds: Jewell, n., dietz, k., farewell, v, 1992.

[26] Andrea Rotnitzky, Quanhong Lei, Mariela Sued, and James M Robins. Improved double-robust estimation in missing data and causal inference models. *Biometrika*, 99(2):439–456, 2012.

[27] Donald B Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976.

[28] Donald B Rubin and Roderick JA Little. Statistical analysis with missing data. *Hoboken, NJ: J Wiley & Sons*, 2002.

[29] Daniel O Scharfstein and Rafael A Irizarry. Generalized additive selection models for the analysis of studies with potentially nonignorable missing outcome data. *Biometrics*, 59(3):601–613, 2003.

[30] Daniel O Scharfstein, Andrea Rotnitzky, and James M Robins. Adjusting for nonignorable drop-out using semiparametric nonresponse models. *Journal of the American Statistical Association*, 94(448):1096–1120, 1999.

[31] Leonard A Stefanski and Dennis D Boos. The calculus of m-estimation. *The American Statistician*, 56(1):29–38, 2002.

[32] Madhukar H Trivedi, James H Kocsis, Michael E Thase, David W Morris, Stephen R Wisniewski, Andrew C Leon, Alan J Gelenberg, Daniel N Klein, George Niederehe, Alan F Schatzberg, et al. Revampresearch evaluating the value of augmenting medication with psychotherapy: rationale and design. *Psychopharmacol Bull*, 41(4):5–33, 2008.

[33] Anastasios A Tsiatis. Improving efficiency and double robustness with coarsened data. *Semiparametric Theory and Missing Data*, pages 221–272, 2006.

[34] Anastasios A Tsiatis, Marie Davidian, and Weihua Cao. Improved doubly robust estimation when data are monotonely coarsened, with application to longitudinal studies with dropout. *Biometrics*, 67(2):536–545, 2011.

[35] Mark J Van der Laan and James M Robins. *Unified methods for censored longitudinal data and causality*. Springer Science & Business Media, 2003.

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

[37] Diane Warden, A John Rush, Madhukar H Trivedi, Maurizio Fava, and Stephen R Wisniewski. The star* d project results: a comprehensive review of findings. *Current psychiatry reports*, 9(6):449–459, 2007.

[38] Stephen R Wisniewski, A John Rush, GK Balasubramani, Madhukar H Trivedi, Andrew A Nierenberg, Stard Investigators, et al. Self-rated global measure of the frequency, intensity, and burden of side effects. *Journal of Psychiatric Practice®*, 12(2):71–79, 2006.

[39] Baqun Zhang, Anastasios A Tsiatis, Eric B Laber, and Marie Davidian. Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*, page ast014, 2013.