Bayesian Methods for SMART Designs With Re-Randomization Based on a Continuous Outcome and Estimating Survival Outcomes When COVID-19 Is a Competing Risk in Cancer Patients

by

Holly Elizabeth Hartman

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Biostatistics)
in the University of Michigan
2021

Doctoral Committee:

Associate Professor Kelley M. Kidwell, Co-chair
Research Professor Matthew J. Schipper, Co-chair
Professor Christopher R. Friese
Associate Professor Rafael Meza
Professor Lu Wang

Holly Elizabeth Hartman
holhart@umich.edu
ORCID iD: 0000-0002-9101-4381

DEDICATION

To my family and friends.
Without your love and support,
I'm not sure where I'd be,
but it surely wouldn't be here.
I love you all.

ACKNOWLEDGMENTS

First, I thank my committee and co-advisors for their guidance. My advisors, Dr. Kelley Kidwell and Dr. Matthew Schipper, have provided immense help and guidance in my academic, professional, and personal development. Dr. Kidwell has provided guidance in both my career development and personal growth as a woman in statistics. I thank her for sharing her insights with me. Dr. Schipper has trained me in how to collaborate and defining work-life balance. From effectively communicating results to setting realistic and healthy deadlines, I thank him for guiding my growth as a statistical collaborator. I'd like to thank the all members of my committee. Thank you Dr. Lu Wang, Dr. Rafael Meza, and Dr. Christopher Friese for your comments and feedback. I'd also like to thank those that have helped with the projects presented here: Dr. Nicholas Henderson, Dan Barker, the Tababoo lab group, the Schipper lab group, and Dr. Roy Tamura. Additionally, thank you to my many classmates and study groups for helping me learn and grow.

I would also like to acknowledge all the other teachers, professors, and mentors that have supported my academic development. There have been many points in my life where I found myself at a crossroads and it is only through the academic guidance I received that I have found myself here. This dissertation could not have been completed without the support, encouragement, and wise words from: Ms. Whitnah, Ms. Anderson, Mr. Morrow, Mr. Richardson, Mr. Paricio, Dr. James Bernard, Dr. Bryan Smith, Dr. Daniel Burgard, Dr. Eric Scharrer, Dr. Leslie McClure, Dr. David Redden, Dr. Daniel Spratt, Dr. Robert Dess, Dr. William Jackson, Dr. Laura Scott, Dr. Jeremy Taylor, Dr. Thomas Braun, and Dr. Bhramar Mukherjee.

Next, I must acknowledge all the incredible support from the administration staff in the Biostatistics office. Thank you for the reminders of deadlines and the assistance navigating the degree process. Your support is invaluable.

My friends and family have also provided invaluable support over the years and this dissertation could not have been completed with them. Thank you Kimberly Dill-McFarland, Justin Leach, Michelle McNulty, Brett Morris, Margaret Banker, Nicole Wakim, Cathy Smith, Erin Loomas, Mom, Dad, Amy, Chris, and Gillian. You have all done a wonderful job dealing with my neuroses and quirks while still making me a better person and researcher through your friendship and love.

Lastly, my husband, Jeremy Frey. There are not enough words for you, my love. Thank you for everything.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE

LIST OF TABLES

TABLE

LIST OF ACRONYMS

ARAMIS A Randomized Multicenter Study for Isolated Skin Vasculitis

CFR case fatality rate

DO dichotomized outcome

DTR dynamic treatment regimen

GEE generalized estimating equations

HCF health care facility

HCW health care worker

HR hazard ratio

IHME Institute or Health Metrics and Evaluation

IPTW inverse probability of treatment weighting

IRDiRC International Rare Diseases Research Consortium

MF mapping function

mpSIR mixing population Susceptible-Infected-Recovered/Removed

NCDB National Cancer Database

NHANES National Health and Nutrition Examination Survey

OWL outcome weighted learning

PRO patient reported outcome

rMSE root mean squared error

RMST restricted mean survival time

SEER Surveillance, Epidemiology, and End Results

SIR Susceptible-Infected-Recovered/Removed

SMART sequential, multiple assignment, randomized trial

snSMART small sample, sequential, multiple assignment, randomized trial

TBRL tree-based reinforcement learning

TF-SMART SMART which uses a tailoring function

TV-SMART traditional SMART using a predefined tailoring variable

TSP treatment specific pathway

WRR weighted and replicated regression

ABSTRACT

Sequential, multiple assignment, randomized trials (SMARTs) typically rely on a binary variable to define response which is used in assigning the next stage treatment assignment. We develop a function of a continuous outcome to assign a probability of staying on the same treatment and then randomly assign the next treatment using a multinomial distribution.

First, we develop a new trial design for small sample SMARTs (snSMARTs). The overall goal of the trial is to determine the optimal first stage treatment. We use a function, called the mapping function, to map the first stage outcome to a probability of staying on the same treatment and Bayesian methods to analyze data from both stages. Re-randomization based on a mapping function of a continuous outcome allows for snSMARTs to be conducted without requiring a binary outcome. We perform simulation studies to compare the proposed design with continuous outcomes to standard snSMART designs with binary outcomes. The proposed design results in more efficient treatment effect estimates and similar outcomes for trial patients. This addresses a gap in rare disease clinical trial methodology by presenting a trial that keeps a continuous outcome continuous, allows patients to either stay or switch treatments based on their outcome, and uses repeated measures for improved statistical properties.

Next, we apply similar concepts to standard size SMARTs with continuous outcomes where the goal is to determine the optimal dynamic treatment regimen (DTR). We present a new trial design for SMARTs that use a tailoring function instead of a binary tailoring variable. Here, we simultaneously develop a tailoring variable and estimate the DTR. We apply methods for developing DTRs from observational data: tree-based regression learning and Q-learning. We compare this to a typical SMART design with a tailoring variable and analyzed with generalized estimating equations. This project addresses a gap in clinical trial methodology by presenting SMARTs where second stage treatment is based on a continuous outcome. We demonstrate that data from a SMART using a tailoring function can be used to efficiently estimate the DTR and is more flexible under varying scenarios than a SMART with a tailoring variable.

Finally, we develop Bayesian methods to estimate overall survival for cancer patients when COVID-19 presents a competing risk. First, we forecast local epidemic trends and daily infection risk for an individual by implementing a truncated Dirichlet state space model. We combine the infection risk with COVID-19 cause-specific mortality to obtain predicted COVID-19 cause-

specific mortality. We then combine the COVID-19 cause-specific mortality as a competing risk with Bayesian, pre-COVID-19 overall survival estimates in cancer patients. Our main outcome of interest is the difference in restricted mean survival time under immediate treatment (7 days after diagnosis) versus delayed treatment (60 days after diagnosis). We present example cases to demonstrate the use of our methods. The goal of our method is to create an overall personalized predicted survival curve with posterior prediction intervals to allow health care providers to make more informed decisions on immediate versus delayed treatment for cancer patients.

CHAPTER 1

Introduction

Dynamic treatment regimens (DTRs) are sequences of treatments that are tailored for each patient based on the evolution of the patient's disease and response to prior treatments. A DTR may take the form of: "Start with treatment A. If X, then proceed to treatment B. Otherwise, proceed to treatment C." This DTR may be written (A, B, C) to indicate the appropriate sequence of treatments. "X" is an intermediate outcome and is called the tailoring variable as it is used to tailor the sequence of treatments that a patient receives. The tailoring variable is an short-term outcome that is observed after the first treatment has been given. DTRs are useful when there are certain patient characteristics that may influence how a person responds to treatment. This allows patients with the same condition to potentially receive different treatments in order to have the greatest improvement.

Sequential, multiple assignment, randomized trials (SMARTs) are a category of trial design that allow researchers to estimate the optimal DTR. SMARTs were described in Murphy (2005) using weighted and replicated regression (WRR) with generalized estimating equations (GEE) to estimate the optimal DTR. In a SMART, patients are randomized to a first stage treatment. After some pre-defined interval, the tailoring variable is measured. The tailoring variable is typically defined as response or non-response to the first stage treatment and must be selected prior to the start of the trial. Then, based on the tailoring variable and design, a second stage treatment is assigned either randomly or deterministically. Once the trial is complete, researchers can use GEE and the replicated and weighting method for analysis (Kidwell et al., 2018; Nahum-Shani et al., 2012a).

SMARTs have since been used in a wide range of fields from oncology (Thall, 2015) to addiction (Lei et al., 2012). For example, the CATIE investigation examined optimal treatment sequences for patients with chronic schizophrenia (McEvoy et al., 2006). In this trial, one outcome of interest was identifying the optimal treatment for patients who have an inadequate response to their first stage treatment. This study found that for patients that didn't improve on their first stage treatment, clozapine was the best next treatment despite the increased risk of serious side effects. Another example is in oncology where patients frequently receive multiple forms of treatment until

they find one that works for their specific cancer. In a prostate cancer study, the overall goal was to identify the best overall regimen of treatments by comparing 12 different two-stage treatment regimens Thall et al., 2007. Both of these studies were designed to better treat diseases where it is common for patients to try multiple treatments. These studies aim to identify the best first stage treatment that maximizes the probability of success, and then the best second stage treatment if the first treatment is not successful.

A similar trial design, called a small sample, sequential, multiple assignment, randomized trial (snSMART), has been developed for rare diseases (Tamura et al., 2016). The progression of the trial is similar, but instead of estimating the optimal DTR, the goal is to estimate the optimal first stage treatment. snSMARTs leverage data from both stages to jointly estimate the best first stage treatment within a Bayesian framework (Wei et al., 2018). snSMARTs also include a dichotomous definition of response to determine the set of possible second stage treatments (Wei et al., 2018).

That is, both SMARTs and snSMARTs require a predefined tailoring variable. There are many issues that arise when using a tailoring variable. First, it can be challenging for researchers to identify what classifies response and non-response, so that defining response could require time consuming and expensive research such as conducting a pilot trial. In many cases, especially in rare diseases, this is not feasible due to lack of funding and available patients. Further, in many cases, the tailoring variable is a dichotomized continuous variable. For example, a tailoring variable may be if a patient's score is above some or below threshold. There are a multitude of outcomes that are not naturally binary that would require developing a dichotomization strategy to be used as a tailoring variable including percent change, patient reported outcomes (PROs), utility measures, and biomarker measurements. Dichotomizing continuous variables can have statistical consequences and result in a loss of information (Altman, 1994; Altman and Royston, 2006; Cox, 2018; Day et al., 2018) which is not trivial in small samples. In the case of snSMARTs, rare disease trials aim to have holistic outcomes that incorporate both clinical features and PROs which may make requiring a binary outcome difficult (Day et al., 2018). In this dissertation, we develop a clinical trial design which does not require the use of a binary tailoring variable to determine the next stage treatment We apply this design in both standard to large sample sizes where interest is in DTRs and in small samples where interest is in pooling data to identify the best first stage treatment.

In Chapter 2, we develop an snSMART design that uses a function to map a continuous outcome to a probability of staying on the same treatment. We call this function a mapping function. In this new snSMART design, patients are randomly assigned to one of three treatments with equal probability. At the end of the first stage, a continuous outcome is measured. Then, based on that outcome and the mapping function, the patient is assigned a probability of staying on the same treatment. Better outcomes map to values closer to 1 so that patients that are doing well on the

treatment have a higher probability of staying on the same treatment. Those that are randomized to switch treatments are then randomized to one of the two remaining treatments with equal probability. After the second stage, the same continuous outcome is measured again. The outcomes from the first and second stage are jointly analyzed using Bayesian models to estimate the first stage treatment effects. We simulated trials to compare this design to a typical snSMART design where the continuous outcome is dichotomized. We compare the accuracy and efficacy of the treatment effect estimates, patient outcomes on the trial, and the percentage of patients on each trial arm. This work has been published in *Statistics in Medicine* (Hartman et al., 2021). Here, we address a gap in rare disease clinical trial methodology by presenting a trial that keeps a continuous outcome continuous, allows patients to either stay or switch treatments based on their outcome, and uses repeated measures for improved statistical properties.

In Chapter 3, we implement a similar trial design as that in Chapter 2, but in larger samples with the overall goal of determining the optimal DTR. Our objective with this SMART design is to define the DTR in terms of the first and second stage treatments and the tailoring variable. This design allows for more flexibility in that it does not require a predefined definition or dichotomous tailoring variable. In order to accomplish this, we implement a generalization of the tailoring variable we term the tailoring function. The tailoring function, like the mapping function from Chapter 2, assigns probabilities of receiving each second stage treatment based on the intermediate outcome. We compare analytic modalities for this design drawing on the methods that have been used in developing DTRs from observational data including Q-learning and tree-based reinforcement learning (TBRL). To compare the designs and analytic methods, we simulate trials under different settings and apply the resulting estimated DTR to a new population of 10,000 patients and compare average second stage outcomes. The goal of this design is to identify the DTR that maximizes the expected second stage outcome. We compare the analytic method (Q-learning vs TBRL) as well as the design (tailoring function vs pre-selected tailoring variable). This project addresses a gap in clinical trial methodology by presenting SMARTs where second stage treatment is based on a continuous outcome.

The next chapter shifts focus from clinical trial design to predicting the optimal timing of cancer treatment in the midst of a pandemic. In areas majorly impacted by the COVID-19 pandemic, health care systems needed new models to prioritize patient care. These measures are needed both to create space for the influx of COVID-19 patients as well as to minimize risk to patients coming in for other types of care. One group of patients particularly impacted by care disruption is cancer patients. Cancer patients may be at increased risk of dying from COVID-19 (Landman et al., 2020; Williams et al., 2020) so it is critical to use epidemiological data of the community surrounding the health care facility (HCF) to determine the patients risk of contracting COVID-19 due to coming to the HCF for cancer treatment. Some forms of treatment require coming

into the HCF multiple times over the course of several weeks or months which may increase the risk of contracting COVID-19 (Center for Disease Control and Prevention, 2020; Nguyen et al., 2020). However, delaying cancer treatment can result in increased cancer mortality (Khorana et al., 2019). In order to help clinicians and patients balance these risks, we developed an online webapp called OncCOVID. This app takes input data regarding the patient's demographics, cancer, and community to output personalized survival curves under both immediate treatment and delayed treatment. This work was also published in *JAMA Oncology* (Hartman et al., 2020). In this chapter, we further develop the methods used in this webapp by applying Bayesian methods to obtain posterior survival estimates and posterior predictive intervals under a range of potential treatment start dates. The methods presented in this chapter introduce an infectious disease as a competing risk to a chronic disease to estimate survival curves and corresponding predictive intervals. We draw on methods across disciplines to estimate survival for cancer patients who either delay their cancer care or proceed with immediate cancer care. The overall goal of this project is to take into account all patient, environment, and disease characteristics that influence cancer and/or COVID-19 survival to output highly personalized survival estimates. Our main outcome of interest is the difference in restricted mean survival time (RMST) defined as the average survival time up until a fixed time-point under immediate treatment (7 days after diagnosis) versus delayed treatment (60 days after diagnosis). These estimates can then be used to guide treatment planning.

CHAPTER 2

Design and Analysis Considerations for Utilizing a Mapping Function in a Small Sample, Sequential, Multiple Assignment, Randomized Trials with Continuous Outcomes

2.1 Introduction

Recent developments have been made in small sample, sequential, multiple assignment, randomized trials (snSMARTs) (Tamura et al., 2016; Wei et al., 2018) where the primary goal is identifying the optimal first stage treatment by sharing information across the two stages of the trial. In snSMARTs, patients are randomly assigned to a first treatment and a binary outcome (e.g. response) is measured at a fixed time point. The second stage treatment is assigned randomly or deterministically based on the design and response to first stage treatment. The same binary outcome is again measured at the end of the second stage. Data from the first and second stages are shared to determine the single optimal first stage treatment. This information-sharing design is beneficial in the setting of rare diseases. The attractiveness of using an snSMART design over a traditional single-stage design is that each patient contributes more information than in a single-stage design. snSMARTs may be preferable to traditional crossover designs in that patients who respond remain on treatment and those who do not respond switch to other treatments. While snSMARTs designs are similar to sequential, multiple assignment, randomized trials (SMARTs) (Murphy, 2005), they differ in their primary goal. Unlike SMARTs, the focus of an snSMART is not on estimating the effects of the embedded dynamic treatment regimens (DTRs) or tailored sequences of treatments, but rather in efficiently using two stages of information from the same individuals to find the optimal first stage treatment.

One ongoing snSMART is the A Randomized Multicenter Study for Isolated Skin Vasculitis (ARAMIS) trial (Micheletti et al., 2020). The goal of this trial is to determine the optimal treatment

for isolated skin vasculitis, a rare disease. Currently, there are three treatments available for patients and all are prescribed routinely, and none are considered standard of care for this disease. As such, this trial has three active treatments as the three arms and there is no control arm. The snSMART design is appropriate since the number of patients a more major limiting factor relative to the duration of the trial. Therefore, the longer duration of the trial due to the two stages was seen as a lower priority than needing to obtain as much information as possible from a limited number of patients.

One benefit of a multistage trial in rare disease settings is that there is more information obtained from an individual patient. The International Rare Diseases Research Consortium (IRDiRC) published a list of recommendations for rare disease clinical trial design in 2018 (Day et al., 2018). In these recommendations, they include using longitudinal data and using patients more than once. Multistage designs have been shown to have increased power relative to single stage designs (Tamura et al., 2011, 2016). This is due to having more information about the sources of variation (Senn, 2016). Single stage designs can only identify between treatment variation (Senn, 2016). Crossover designs and other multistage designs also can identify the between-patient variation (Senn, 2016) resulting in less variation in the error term and more statistical efficiency. Even more powerful are snSMARTs where patients may stay on the same treatment so the variation between treatments and the variation between patients receiving the same treatment (Senn, 2016) can both be identified.

Multistage designs such as crossover designs do not incorporate the patient outcome in the design. This is not ideal for rare diseases (Hilgers et al., 2018; Mullins et al., 2014). Second stage treatment assignment in snSMARTs and other multistage designs is often based on a binary outcome such that the trial protocol must specify a binary rule for what qualifies as a "response." Some procedures may dichotomize a continuous variable such as change in prostate specific antigen (PSA) in a prostate cancer trial (Thall, 2015) or response may be determined using defined criteria such as Response Evaluation Criteria In Solid Tumors (RECIST). In an snSMART, the binary outcome must be the same at the end of the first and second stages and must be selected and agreed upon before the start of the trial (Lei et al., 2012).

However, in many studies the outcome of interest may not be a binary variable or there may not be enough information known about a continuous variable to dichotomize it for use in second stage randomization. Another recommendation from IRDiRC is to not dichotomize continuous endpoints (Day et al., 2018). Dichotomizing a continuous outcome often results in lowered statistical efficiency (Altman, 1994; Altman and Royston, 2006; Cox, 2018; Day et al., 2018). In rare diseases, ensuring statistical efficiency is critical due to the inherently low number of patients in the trial (Cox, 2018). Further, results may vary based on the dichotomization strategy and cut off selected (Schellingerhout et al., 2009).

In rare diseases and rare cancers, selecting a robust, holistic outcome is ideal to capture the patient's subjective experience and individual variations in the disease (Cox, 2018; Day et al., 2018; Tudur Smith et al., 2014). Patient reported outcomes (PROs) can be used to measure treatment benefit or risk (US Department of Health and Human Services Food and Drug Administration, 2009) and can record the full patient experience with a treatment. The primary endpoint of the ARAMIS trial is a dichotomized composite endpoint that incorporates PROs and clinical outcomes (Micheletti et al., 2020). The goal of the composite endpoint is to holistically capture both the disease clinical progression and the patient experience on the treatment. Dichotomized outcomes may miss critical information and may obfuscate results. A major drawback of implementing an snSMART in oncology and other chronic diseases is that efficacy and toxicity are both important outcomes and creating a dichotomous binary outcome that captures both outcomes effectively can be challenging (Thall, 2015). Dichotomizing a continuous outcome, such as a PRO or other composite endpoint, requires expansive knowledge on the outcome and treatment effects in order to select an appropriate dichotomization strategy.

Moreover, when the outcome of interest is a continuous variable, a clear choice for a dichotomization method or binary surrogate may not always be available prior to the start of a study. If a continuous variable is dichotomized, the cut off selected will play a critical role in the trial progression. If the cut off selected is too high and the majority of participants are categorized as non-responders, then most patients will switch treatments and the long term benefits of a treatment may not be observed. Conversely, if a cut off is selected that is too low and the majority of participants are categorized as responders, then most patients will stay on the treatment and the effects of switching treatments will not be adequately observed and estimation errors in treatment effects may occur.

In small samples, this problem of identifying an appropriate cut off is magnified as fewer patients in each treatment pathway will be observed with poorly selected cut offs. Further, if a dichotomous response is used for the endpoint or outcome of a trial, as in a responder analysis, there is potential for significant loss in statistical power (Cox, 2018; Snapinn and Jiang, 2007). In small samples and rare disease settings, there may be few prior studies and less knowledge regarding the treatment effects on the outcome which further hinders the ability to select a suitable binary outcome. Additionally, the time and resources to run a pilot study to determine an appropriate cut off are often not available in diseases and disorders that affect a small number of individuals.

Here we examine an snSMART design with a continuous outcome measured at the end of stages one and two where re-randomization depends on a mapping function of a continuous outcome as opposed to a binary outcome. This design leverages the additional information gained by a multistage design and better captures the effect of the treatment by using a continuous outcome. The ultimate aim is to estimate the expected first stage treatment effects for the multiple treatments be-

ing investigated in the snSMART. As in the ARAMIS trial, we assume that all treatments are used in practice and the goal is to identify the optimal first stage treatment. Examples of continuous outcomes with no clear dichotomization strategy are percent change, PROs, utility measures (or some combination of efficacy and toxicity), probabilities of response, or other composite endpoints. The trial design presented here allows for the implementation of an snSMART without requiring a binary first stage outcome to determine the next stage treatment for a patient. This method maintains the patient benefit of an snSMART by having an increased probability of switching treatments if an individual is not responding well on the current treatment and increased probability of staying on that treatment if the individual is responding well. It also maintains the benefit of being able to identify variation between treatments and between patients resulting in a more powerful design. In addition to allowing randomization to depend on a continuous outcome in the snSMART design, the methods presented here allow for continuous first and second stage outcomes as opposed to previous methods to analyze snSMART data that only applied to binary outcomes. Thus, due to the use of continuous outcomes, we cannot use previously proposed methods for snSMARTs (Wei et al., 2018) and due to on second stage randomization depending on a continuous outcome and sharing information across stages for the first stage treatment effect, we cannot use standard SMART methods.

2.2 Methods

2.2.1 snSMART design with a continuous mapping function

Here we present the proposed snSMART design which allows for a continuous outcome measured at the end of the first and second stages. This design does not require a binary outcome. In this design, all patients are randomized with equal probability to a first stage treatment. We measure the continuous outcome at the end of stage one for patient i, Y_{i1}, at a set time point, t. The probability of staying on the same treatment for patient i is a function of Y_{i1}, and we call this function, $f(Y_1)$, the mapping function. The mapping function ranges between 0 to 1 in order to yield valid probabilities. This is similar to the idea of treatment effect mappings introduced by Rosenberger (1993), but only relies on the patient's individual outcome and not the outcome of others. Patients randomized to switch treatments are randomized to one of the remaining treatments with equal probability. At a set final time point, $2t$, the outcome at the end of stage two, Y_{i2}, is measured. A schematic of the traditional snSMART design is given in Figure 2.1A and the proposed snSMART design is presented in Figure 2.1B.

The randomization at the end of the first stage can be thought equivalently as a one-step randomization process or a two-step randomization process. The one step randomization proceeds as

presented in Figure 2.1B where the next treatment is decided by a multinomial distribution with probabilities $[f(Y_1), 0.5(1 - f(Y_1)), 0.5(1 - f(Y_1))]$ for staying on the same treatment and switching to either other treatment, respectively. It can also be thought of as a two-step randomization process with two binomial distributions (Figure 2.2). The first randomization determines if the patient stays or switches treatment and has probability of staying equal to $f(Y_1)$. If the person switches treatments, then there is a second randomization with probability 0.5 for each remaining treatment.

2.2.1.1 The mapping function

The mapping function maps the first stage outcome, Y_{i1}, to $[0, 1]$ and gives the probability of staying on the same treatment. More favorable values of the outcome map to values closer to 1 so that patients doing well on a treatment have a higher probability of staying on the same treatment. Depending on what is known about the outcome, various mapping functions can be selected. We provide examples of mapping functions and assume that higher values of Y indicate more favorable performance, but the functions are generalizable in other cases. While we offer the following suggestions, any cumulative distribution function can be used.

For areas of study with small samples such as rare diseases, there may be little information known about the distribution of Y_1. If only minimum and maximum values are known for Y_1, a linear mapping function $f(Y_1) = (Y_1 - Y_{min})/(Y_{max} - Y_{min})$ can be used. This linear $f(Y_1)$ function can be easily modified using powers for different trial characteristics. For example, if it is desirable for high proportions of patients stay on the same treatment through both stages, or if it is expected that the outcome may be right skewed, then using $f(Y_1)^{1/k}, k > 1$, may be appropriate. Likewise, if having more people switch is of interest or if the outcome is left skewed, then $f(Y_1)^k, k > 1$, may be appropriate.

If there are extreme values within this possible range of Y_1 that rarely occur, the linear function over the possible range of Y_1 may result in very few people on certain treatment regimens. For example, if Y_1 has a distribution as seen in Figure 2.3 (iii), then a linear mapping function will result in the majority of people switching treatments. To adjust for this, one can select a more practical minimum and maximum and then truncate the mapping function at 0 and 1. These values may represent safety and ethical limits in addition to the practicality of the study. The minimum could be the lowest value of Y_1 where investigators feel comfortable with a patient staying on the same treatment. This lower lower limit is a safety measure as it represents the worst outcome a patient will have and stay on the same treatment. The highest value of Y_1 is selected where investigators feel comfortable with a patient switching treatments, so that if a patient has a higher outcome they will stay on the current treatment. This upper limit represents an ethical boundary since it could be viewed as unethical to switch treatments for a patient responding this well to

Figure 2.1: A traditional snSMART design with a binary outcome (A) and the proposed snSMART design with a continuous outcome (B). Each stage is the same length so that Y_1 is measured at time t and Y_2 is measured at time $2t$. A, B, and C are treatments, R indicates a randomization point, $f(Y_1)$ is the mapping function, expressions indicate probability of assignment to the treatment following, Y_1 is the first stage outcome, and Y_2 is the second stage outcome.

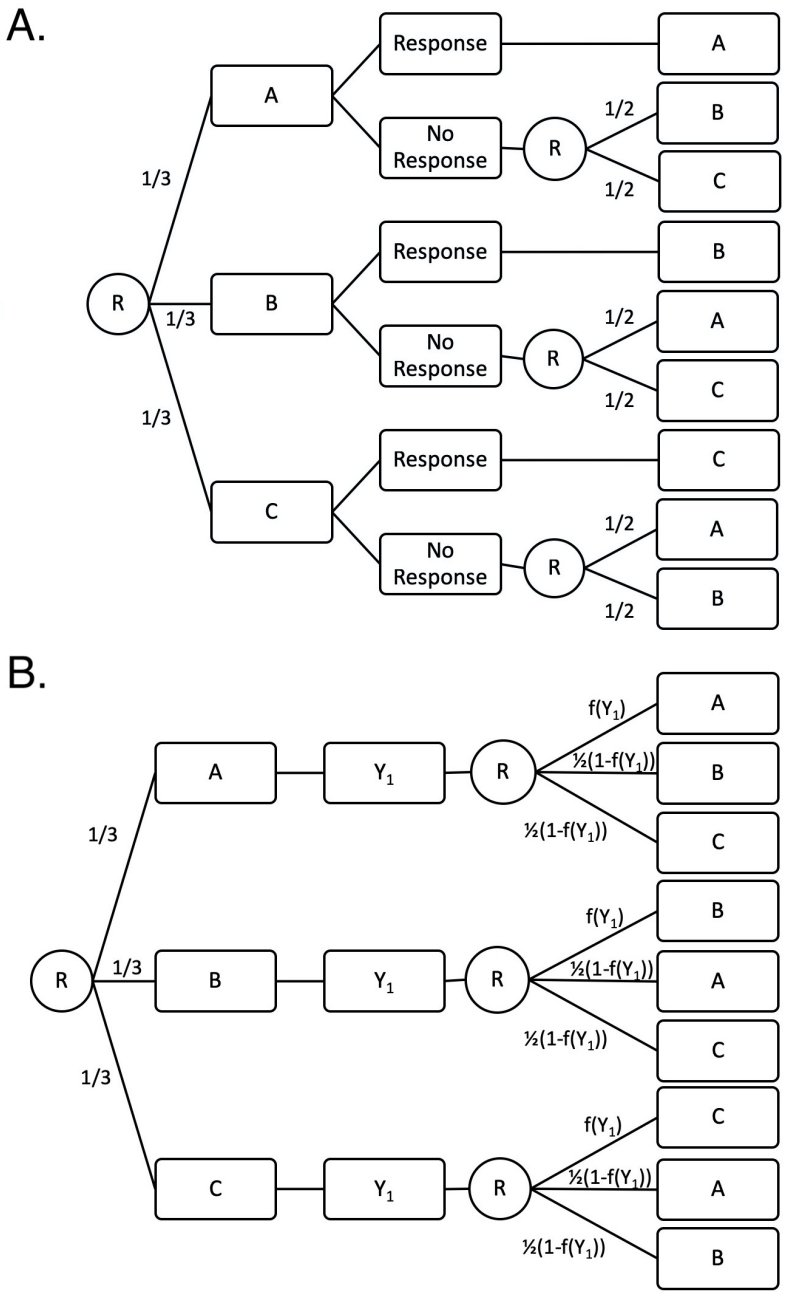

Figure 2.2: The proposed snSMART design with a mapping function and two-step randomization. Each stage is the same length so that Y_1 is measured at time t and Y_2 is measured at time $2t$. A, B, and C are treatments, R indicates a randomization point, $f(Y_1)$ is the mapping function, expressions indicate probability of assignment to the treatment following, Y_1 is the first stage outcome, and Y_2 is the second stage outcome.

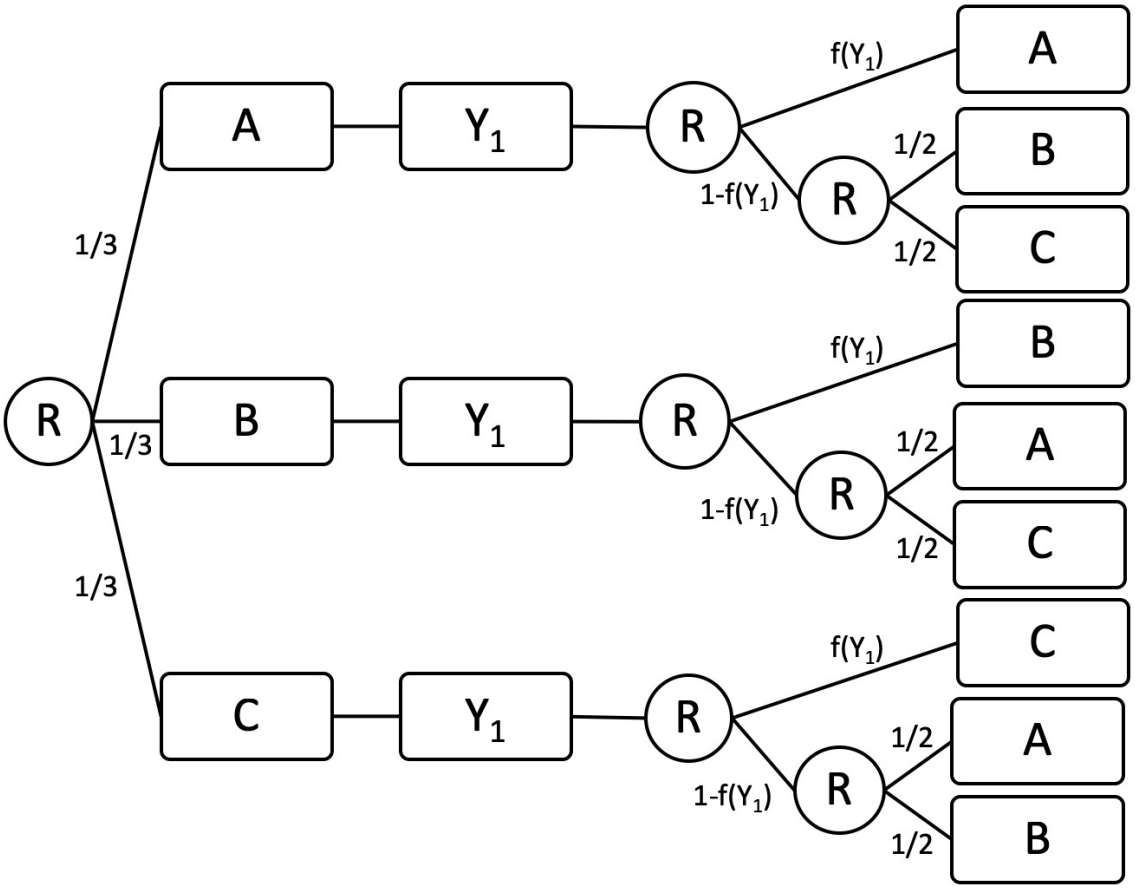

their current treatment. Patients can be consulted to aid in determining the upper and lower limits, which may make the trial more patient centered and appealing to patients (Gaasterland et al., 2019). Developing these extreme limits requires less prior knowledge regarding the treatment effects than is needed to effectively dichotomize a continuous outcome. Additionally, even when minimums and maximums are selected, the mapping function is more flexible than a binary outcome.

Figure 2.3: Possible distributions of outcome Y and mapping functions that balance patients staying or switching treatments. The plots labelled "Distribution of Y" are the probability distribution functions. The y-axis for plots labelled "Possible mapping functions" is the probability of staying on the treatment for a given value of Y. Higher values of Y are assumed to indicate better outcomes.

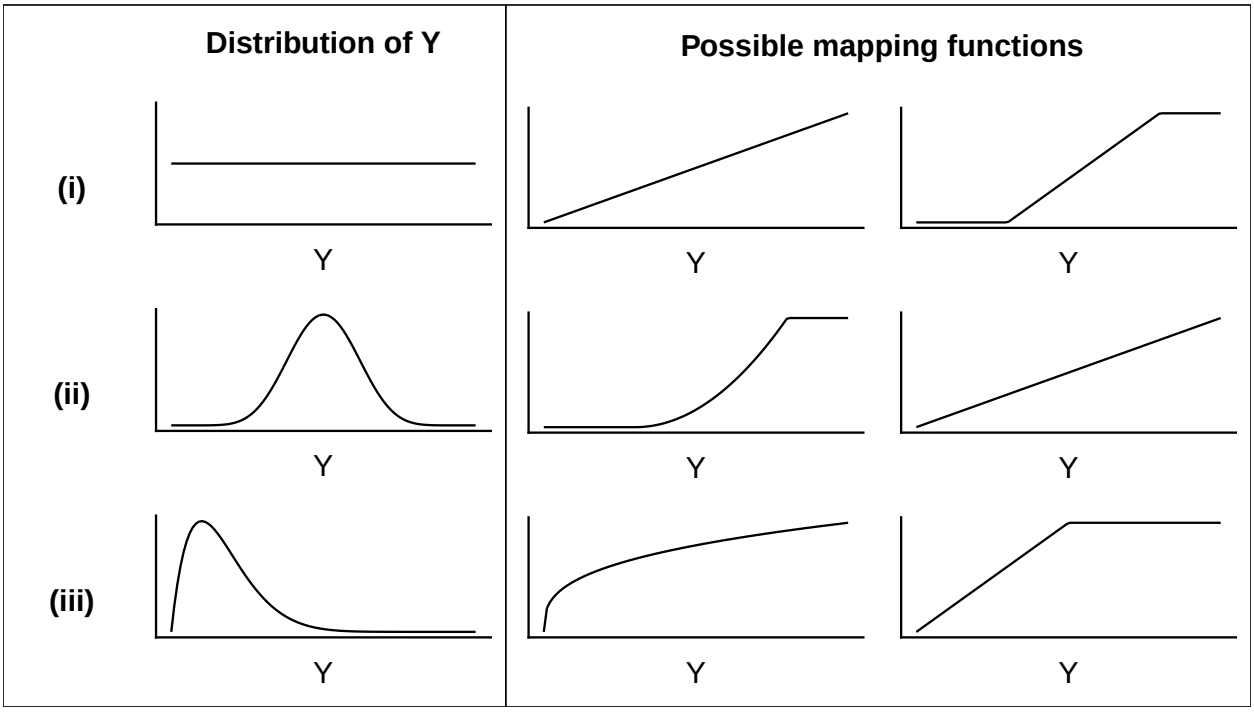

2.2.2 Bayesian analysis methods

2.2.2.1 Model and likelihood

We assume the data from the snSMART design, has a multivariate normal likelihood (Zajonc, 2012):

$$\begin{bmatrix} Y_{i1} \\ Y_{i2} \end{bmatrix} | T_{i1}, T_{i2} \sim MVN \left(\begin{bmatrix} \mu_1(T_{i1}) \\ \mu_2(T_{i1}, T_{i2}) \end{bmatrix}, \mathbf{V}(T_{i1}, T_{i,2}) \right)$$

where Y_{is} is a continuous outcome and T_{is} is the treatment for patient i in stage s, $s = 1, 2$. The mean outcome for stage one, μ_1, is a function of only the first stage treatment. The mean for the

second stage, μ_2, is a function of both the first and second stage treatments. The covariance matrix is a function of the sequence of treatments the patient received.

The mean treatment effects for stage one and two are modeled as follows for treatments A, B, and C:

$$\mu_1(T_{i1}) = \sum_{j=A}^{C} \beta_j I(T_{i1} = j) \tag{2.1}$$

$$\mu_2(T_{i1}, T_{i2}) = \alpha_1 \sum_{j=A}^{C} \beta_j I(T_{i1} = j) + \alpha_2 \sum_{k=A}^{C} \beta_k I(T_{i2} = k) + \alpha_3 I(T_{i1} = T_{i2})$$

The β_j parameters are the expected effect of treatment j, $j = A, B, C$, in the first stage. The primary goal of an snSMART is to estimate these β_j parameters in order to identify the single optimal treatment during the first stage. The first stage mean outcome depends on the treatment effect for the treatment received in stage one. We model the second stage mean outcome as a weighted average of the treatment effects from stage one and stage two with an additional effect if the patient stays on the same treatment. Practically, this mean model for stage two can occur when there is some lingering effect of the first treatment (α_1) and some additional effect of the second treatment (α_2) when $\alpha_1, \alpha_2 > 0$. Alternatively, if the patient stays on the same treatment, the effect of the treatment in the second stage is the first stage effect with some cumulative effect that occurs on the treatment longer term (α_3). The proposed mean models allow for shared information between the two stages through β_j. The previously proposed linkage parameters in prior work to analyze data from an snSMART with binary outcomes (Wei et al., 2018) are similar to the α parameters that we use in our models.

For modeling the covariance, we use $\mathbf{V}(T_{i1}, T_{i,2}) = V_1 I(T_{i1} = T_{i2}) + V_2 I(T_{i1} \neq T_{i2})$ where V_1 and V_2 are both 2×2 variance-covariance matrices. This allows those who stay on the same treatment to have a different correlation between stage one and stage two outcomes than those who switch treatments.

Data from the snSMART design in Figure 2.1B can be analyzed using other mean models (e.g. to consider non-additive effects or effects of treatment specific pathways), but since our focus is on small samples, we present a parsimonious model where information sharing is feasible and the primary goal is to estimate and compare first stage treatments.

2.2.2.2 Prior distributions

We propose minimally informative prior distributions for all model parameters. For the treatment effects, βs, we use a normal prior with mean within the range of outcome Y. We suggest using the midpoint between Y_{min} and Y_{max} for the mean.

We impose three constraints regarding the α_m parameters, $m = 1, 2, 3$:

Constraint 1: $\alpha_2 = 1 - \alpha_1$, $\alpha_1, \alpha_2 > 0$,

Constraint 2: $\alpha_2 > \alpha_1$, and

Constraint 3: $\alpha_3 \geq 0$.

Constraint 1 states that if a person does not stay on the same treatment then the second stage expected outcome is a weighted average of expected first stage outcomes of the stage one and stage two treatments. This constraint facilitates estimation by requiring fewer variables to be estimated. Constraint 2 states that the second stage treatment has a larger effect on the second stage outcome than the first stage treatment. Constraint 3 states that staying on the same treatment has a non-negative effect on the expected second stage outcome. Constraints 2 and 3 reduce the parameter space and are therefore practical in many rare disease settings.

The prior for α_1 is set to be uniform from 0 to 0.5. Since we set $\alpha_2 = 1 - \alpha_1$ per constraint 1, this prior fulfills constraint 2 and forces $\alpha_2 > \alpha_1$ since α_1 will always be less than 0.5. Since we constrain α_3 to be non-negative, we use a folded normal (FN) distribution. For the prior on the covariance matrices, we use inverse Wishart (IW) distributions since they are the non-informative conjugate prior for multivariate normal covariances. Thus, suggested priors are as follows: $\beta_j \sim N(\text{mean} = 50, \text{standard deviation (sd) = } 50)$ for all j, $\alpha_1 \sim Unif(0, 0.5)$, $\alpha_3 \sim FN(\text{mean = } 0, \text{sd} = 20)$, and V_1 and $V_2 \sim IW_2\left(\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, 2\right)$.

2.2.2.3 Single-stage design

To examine the efficiency of the trial design of an snSMART with a mapping function, we compare it to a traditional single-stage design in which individuals are equally randomized to one of the three treatments. In the single-stage design, we analyze the outcome using the following Bayesian methods. Using the same notation as in Section 2.2.2, we assume a normal likelihood:

$$Y_{i1}|T_{i1} \sim N\left(\mu_1(T_{i1}), \sigma^2\right)$$

The mean $\mu_1(T_{i1})$ is the same as in (2.1) and the βs have the same priors as in Section 2.2.2.2. The prior for σ is an inverse gamma, $IG(1000, 1000)$.

2.3 Simulations

For each scenario, simulations were done with $n = 100, 30, 15$ in each treatment arm for total sample sizes of 300, 90, or 45. We focus primarily on the results for $n = 30$. Results are based on

2500 simulated trials. For analysis, we used the Bayesian analysis methods described in Section 2.2.2 with 1000 burn in and 5000 MCMC repetitions. We used R and rjags for simulations.

Our primary goal is accurate and efficient estimation of the effects of the treatments after the first stage in order to identify the optimal overall treatment. As such, we examine bias and root mean squared error (rMSE) of the estimates of β_j. Additionally, we examined patient outcomes and how many patients stayed on their current treatment or switched treatments.

2.3.1 Generation of trial data

Trial data was generated by first assigning n patients equally to each of three treatments. Stage one outcome was generated as $Y_1 \sim N(\mu_1(T_1), \sigma^2)$ where σ^2 was the marginal variance for the first stage outcome. Based on this outcome, a patient was randomly assigned to a stay on the same treatment with probability $f(Y_1)$ using $Bernoulli(f(Y_1))$ where $f(\cdot)$ was the mapping function. We set the parameters so that the range of Y was approximately 0-100. Correspondingly we set $Y_{min} = 0$ and $Y_{max} = 100$. We then investigated the three following mapping functions:

$$\begin{aligned} \text{MF 1} &= f_1(Y_1) = Y_1/100 \\ \text{MF 1/2} &= f_2(Y_1) = (Y_1/100)^{1/2} \\ \text{MF 2} &= f_3(Y_1) = (Y_1/100)^2 \end{aligned} \tag{2.2}$$

For Y_1 outside the range of 0-100, these probabilities were truncated to be 0 or 1 respectively. Guidance for selecting the mapping function is provided in Section 2.2.1.1. If the patient did not stay on the same treatment, then the second treatment was randomly assigned using $Bernoulli(0.5)$.

Let the ith row and jth column in $\mathbf{V}(T_{i1}, T_{i,2})$ be denoted $\mathbf{V}_{ij}$. Then the conditional means for the outcomes are:

$$\begin{aligned} Y_{i1}|T_{i1}, T_{i2} &\sim N(\mu_1(T_{i1}), \mathbf{V}_{11}) \\ Y_{i2}|Y_{i1}, T_{i1}, T_{i2} &\sim N\left(\mu_2(T_{i1}, T_{i2}) + \frac{\sqrt{\mathbf{V}_{22}}}{\sqrt{\mathbf{V}_{11}}}\mathbf{V}_{12}(Y_{i1} - \mu_1(T_{i1})), (1 - \mathbf{V}_{12}^2)\mathbf{V}_{22}\right) \end{aligned} \tag{2.3}$$

The second outcome was generated from a conditional normal distribution in equation 2.3. For the single-stage design, only the first stage outcomes were used.

2.3.2 Scenarios

2.3.2.1 Ideal situations

As in Section 2.2.2.1, we let the covariance matrix be $\mathbf{V}(T_{i1}, T_{i,2}) = V_1 I(T_{i1} = T_{i2}) + V_2 I(T_{i1} \neq T_{i2})$ and we set V_1 and V_2 to be:

$$V_1 = \sigma^2 \begin{bmatrix} 1 & \tau_1 \\ \tau_1 & 1 \end{bmatrix}, V_2 = \sigma^2 \begin{bmatrix} 1 & \tau_2 \\ \tau_2 & 1 \end{bmatrix}$$

We set $\tau_1 = 0.8$, $\tau_2 = 0.3$, and $\sigma = 20$ (based on results from a validated ANCA-associated vasculitis PRO present in Robson et al. (2018)). For the mean model, we set $\alpha_1 = 0.2$, $\alpha_3 = 0.8$, $\alpha_3 = 5$. We changed the β parameters to examine different effects of the mapping functions in scenarios such as presented in Figure 2.3. The β parameters are presented in Table 2.1 where Scenarios 1, 2, and 3 are ideal scenarios where the model assumptions are met. These β values reflect possible values for a continuous outcome that has been standardized to range from 0-100 (typical in PRO development) (Robson et al., 2018; Wei et al., 2000).

Table 2.1: Simulation scenarios. Parameters are estimated as in equation 2.1 with any violations as described in Section 2.3.2.3. TSP violates the mean model assumptions, Variance indicates that there are more variance parameters than are estimated, and Correlation indicates that there are more correlation parameters than estimated. β_j indicates the treatment effect of treatment j after a fixed amount of time and the study goal is to accurately and efficiently estimate these parameters. TSP = treatment specific pathway

	β_j			Violation in Assumptions		
Scenario	A	B	C	treatment specific pathway (TSP)	Variance	Correlation
1	40	50	60			
2	20	30	40			
3	60	70	80			
4	40	50	60	×		
5	40	50	60		×	
6	40	50	60			×
7	40	50	60	×	×	
8	40	50	60	×		×
9	40	50	60		×	×
10	40	50	60	×	×	×

2.3.2.2 Comparison with dichotomized outcome

This trial design that does not require a binary outcome (Figure 2.1B) will be primarily useful in cases when a binary outcome is not available. However, in simulations we are able to compare the

proposed design with a mapping function to an snSMART with a binary outcome (Figure 2.1A), even though in practice, a design with a well selected binary outcome may not be always feasible. We compare the three mapping functions in equation 2.2 with three binary outcomes based on dichotomizing the continuous outcome at the end of the first stage. If the patient's outcome at the end of the first stage is above the cut off, they stay on treatment, and if it is below the cut off, they are equally randomized to one of the two remaining treatments. We investigated three cut offs of 30, 50, and 70 for the binary outcome. We use Scenarios 1, 2, and 3 (scenarios with no model assumption violations) in table 2.1 with these cut offs to examine trial properties and compare with the mapping functions in equation 2.2. We compare the average patient outcomes and the number of patients on each sequence of treatments between the trial designs.

2.3.2.3 Model assumption violations

We examined three potential assumption violations individually and in combination as defined in Table 2.1, Scenarios 4-10. In each scenario, the parts of the data generative model without assumption violations are the same as in Scenario 1. First, we assumed the second stage mean was fully dependent on the combination and ordering of treatments the patients received in stages one and two (T_{i1}, T_{i2}) (i.e. not a weighted mean). We call this the treatment specific pathway (TSP) (column "TSP" in Table 2.1; Scenarios 4, 7, 8, 10) and is essentially where there is an interaction between treatments. We set $\mu_2 = \sum_j \sum_k \mathbf{T}_{jk} I(T_{i1} = j, T_{i2} = k)$ where

$$\mathbf{T} = \begin{bmatrix} 55 & 43 & 47 \\ 40 & 50 & 60 \\ 70 & 65 & 60 \end{bmatrix}$$

and $\mathbf{T}_{jk}$ indicates the jth row and kth column of matrix $\mathbf{T}_{jk}$.

The second assumption violation was that we set the standard deviation of the outcome, σ, to depend on treatment (column "Variance" in Table 2.1; Scenarios 5, 7, 9, 10). Here we set the standard deviation to be 10, 20, and 30 for treatments A, B, and C respectively.

The third assumption violation allowed the correlation between the first and second stage outcome to depend on the TSP (column "Correlation" in Table 2.1; Scenarios 6, 8-10). Practically, this would arise if treatment outcomes have different correlation due to similarities in treatment mechanism even if the actual treatments differ. We set the correlations $\tau = \sum_j \sum_k \mathbf{R}_{jk} I(T_{i1} = j, T_{i2} = k)$ where

$$\mathbf{R} = \begin{bmatrix} 0.85 & 0.6 & 0.2 \\ 0.6 & 0.8 & 0.5 \\ 0.2 & 0.2 & 0.9 \end{bmatrix}$$

and $\mathbf{R}_{jk}$ indicates the jth row and kth column of matrix $\mathbf{R}_{jk}$.

To assess sensitivity to normally distributed outcomes, we also examined scenarios where the first stage outcome was a scaled Beta distribution rather than a normal distribution. The second stage was still conditionally normal but no longer marginally normal. We set the parameters of the Beta distribution such that the scaled means were the same as those in Scenarios 1, 2, and 3. We set the variance to be 0.1 prior to scaling.

2.4 Results

2.4.1 Estimation of the treatment effect

2.4.1.1 Two-stage versus single-stage design

We first compare this two-stage snSMART design to a single-stage design with the same total number of patients but only one stage of treatment (and thus one outcome per patient). As expected, in Scenarios 1, 2, and 3 we found large reductions in rMSE when comparing the two-stage design with the single-stage design regardless of the mapping function used (Table 2.2). The percent of trials that correctly identified the optimal treatment (determined by having the highest estimated first stage treatment effect) was consistently higher in the two-stage designs than the single-stage designs (Table 2.2), although all simulations resulted in $>95\%$ of trials identifying the true optimal treatment.

2.4.1.2 Mapping function versus dichotomized outcome

In ideal scenarios (1-3), where the data generation model matched the analysis model, our method estimated the first stage treatment effects, β_j with minimal bias (Figure 2.4A). Bias and efficiency did not change substantially by changing the mapping function in Scenarios 1, 2, and 3. Since the main goal of this trial design is to identify the optimal treatment from estimating the β_j parameters, low bias and high efficiency are both desirable trial traits.

Efficiency was similar between the mapping functions for each scenario (data not shown). Correspondingly, the coverage probability of the credible interval was close to the desired 95% and the credible interval width was consistent (Table 2.3). The coverage probability was lower for the small sample size of 15 patients per treatment arm (Table 2.4). Overall, the selection of the mapping function does not appear to have a substantial impact on the operating characteristics of the trial design when assumptions are met.

The results from using a dichotomized outcome vary considerably (Figure 2.4A). When the median treatment effect was selected as the cut off, the bias was lowest, but the efficiency was worst

Table 2.2: The rMSE for each $\beta_j, j = A, B, C$ when estimated in a single-stage design and the proposed two-stage snSMART design. β_j indicates the treatment effect. Ratio is the rMSE divided by the rMSE of the single-stage design for that scenario. rMSE = root mean squared error, MF = mapping function, DO = dichotomized outcome (Section 2.3.1).

		rMSE			Ratio			
Scenario	Design	β_A	β_B	β_C	β_A	β_B	β_C	Percent Correct
1	Single-Stage	3.65	3.69	3.51	1.00	1.00	1.00	97.40
	MF 1	2.97	3.01	2.87	0.82	0.82	0.80	99.52
	MF $1/2$	3.02	2.93	2.84	0.83	0.79	0.81	99.40
	MF 2	3.00	3.09	3.15	0.82	0.84	0.90	99.24
	DO 50	5.57	5.58	5.15	1.53	1.51	1.47	99.16
	DO 30	5.32	3.92	3.19	1.46	1.06	0.91	99.35
	DO 70	2.87	3.52	4.80	0.79	0.95	1.37	99.15
2	Single-Stage	3.51	3.66	3.66	1.00	1.00	1.00	96.64
	MF 1	3.08	2.96	2.91	0.88	0.81	0.79	99.68
	MF $1/2$	3.22	2.99	2.91	0.92	0.82	0.79	99.52
	MF 2	2.94	2.87	2.97	0.84	0.78	0.81	99.44
	DO 50	2.95	3.51	4.79	0.84	0.96	1.31	99.68
	DO 30	5.75	5.50	5.03	1.64	1.51	1.37	99.32
	DO 70	2.98	3.11	3.31	0.85	0.85	0.91	99.64
3	Single-Stage	3.64	3.66	3.67	1.00	1.00	1.00	97.36
	MF 1	3.04	3.00	2.92	0.83	0.82	0.79	99.28
	MF $1/2$	3.02	2.90	2.87	0.83	0.79	0.78	99.60
	MF 2	3.19	3.22	3.19	0.88	0.88	0.87	99.08
	DO 50	5.20	3.81	3.15	1.43	1.04	0.86	98.67
	DO 30	2.99	2.96	2.89	0.82	0.81	0.79	99.40
	DO 70	5.65	5.74	5.23	1.55	1.57	1.43	96.87

Figure 2.4: A. Box plots of bias averaged across the three $\beta_j, j = A, B, C$ parameters for ideal scenarios where β_j indicates the treatment effect. A single point represent the average bias for one simulated trial. Light gray boxes are mapping functions (MF, Section 2.3.1) and dark gray boxes are dichotomized outcomes (DO, Section 2.3.2.2). B. Box plots of bias for each of the three β_j parameters in scenarios where assumptions are violated using MF 1. A single point represent the bias for one β_j in one simulated trial. Top labels indicate which scenario is being presented. Shades indicate the different β_j parameters.

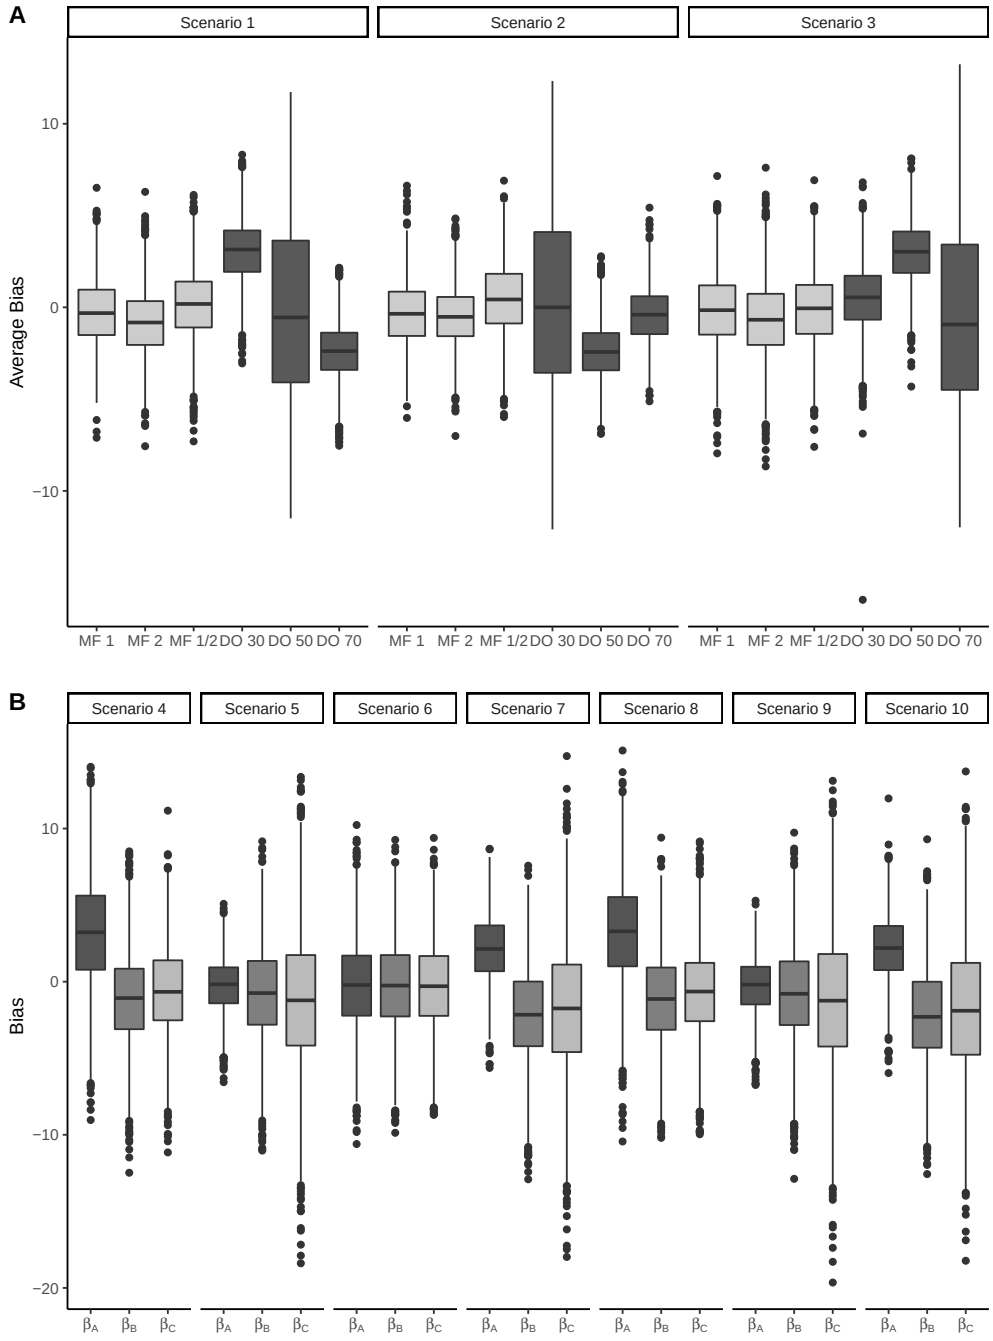

Table 2.3: Coverage probability (CP) and credible interval width (CW) for the ideal Scenarios 1-3 (Section 2.3.2.1) where $n = 30$ for mapping functions and dichotomized outcomes. CP should be close to 0.95 and for CW smaller is better. MF = mapping function, DO = dichotomized outcome, $\beta_j, j = A, B, C$ parameters indicate treatment effects.

CP								
		MF				DO		
Scenario	MF	β_A	β_B	β_C	DO	β_A	β_B	β_C
1	MF 1	0.949	0.935	0.949	DO 50	0.816	0.772	0.820
	MF 1/2	0.944	0.950	0.949	DO 30	0.675	0.811	0.875
	MF 2	0.939	0.934	0.923	DO 70	0.916	0.857	0.698
2	MF 1	0.933	0.942	0.947	DO 50	0.914	0.858	0.696
	MF 1/2	0.930	0.945	0.940	DO 30	0.820	0.793	0.825
	MF 2	0.939	0.939	0.934	DO 70	0.926	0.915	0.898
3	MF 1	0.943	0.944	0.949	DO 50	0.687	0.822	0.888
	MF 1/2	0.941	0.948	0.941	DO 30	0.947	0.934	0.937
	MF 2	0.933	0.933	0.935	DO 70	0.803	0.760	0.810
CW								
		MF				DO		
Scenario	MF	β_A	β_B	β_C	DO	β_A	β_B	β_C
1	MF 1	11.560	11.358	11.239	DO 50	14.926	14.655	14.050
	MF 1/2	11.561	11.263	11.071	DO 30	11.605	10.376	9.791
	MF 2	11.192	11.187	11.260	DO 70	10.223	10.342	10.891
2	MF 1	11.395	11.221	11.164	DO 50	10.213	10.374	10.910
	MF 1/2	11.840	11.356	11.091	DO 30	15.223	14.854	14.086
	MF 2	11.090	10.959	11.049	DO 70	10.939	10.767	10.952
3	MF 1	11.783	11.554	11.373	DO 50	11.650	10.431	9.854
	MF 1/2	11.483	11.334	11.194	DO 30	11.172	10.849	10.739
	MF 2	11.863	11.837	11.727	DO 70	14.795	14.519	14.049

Table 2.4: Coverage probability (CP) and credible interval width (CW) for ideal Scenarios 1-3 (Section 2.3.2.1), n per arm $= 100, 15$. MF = mapping function, $\beta_j, j = A, B, C$ parameters indicate treatment effects.

			CP			CW		
n	Scenario	MF	β_A	β_B	β_C	β_A	β_B	β_C
100	1	MF 1	0.949	0.943	0.951	6.511	6.393	6.387
		MF 1/2	0.936	0.934	0.941	6.514	6.350	6.261
		MF 2	0.929	0.918	0.907	6.228	6.185	6.368
	2	MF 1	0.940	0.939	0.935	6.346	6.230	6.302
		MF 1/2	0.923	0.933	0.935	6.681	6.401	6.288
		MF 2	0.951	0.930	0.908	6.159	6.032	6.234
	3	MF 1	0.939	0.934	0.939	6.706	6.559	6.454
		MF 1/2	0.943	0.932	0.936	6.516	6.406	6.323
		MF 2	0.947	0.937	0.943	6.757	6.707	6.738
15	1	MF 1	0.931	0.922	0.935	15.928	15.647	15.402
		MF 1/2	0.938	0.937	0.934	15.878	15.539	15.299
		MF 2	0.929	0.916	0.914	15.566	15.394	15.395
	2	MF 1	0.931	0.919	0.930	15.821	15.523	15.343
		MF 1/2	0.934	0.929	0.941	16.263	15.619	15.239
		MF 2	0.905	0.892	0.893	15.434	15.218	14.992
	3	MF 1	0.929	0.932	0.924	16.112	15.848	15.598
		MF 1/2	0.913	0.920	0.913	15.787	15.669	15.467
		MF 2	0.921	0.916	0.922	16.224	16.010	15.935

as seen by the high variance (Figure 2.4A). Moderate bias was observed for other cut offs due to poor estimates of the correlation parameters and $\alpha_m, m = 1$ and 3. The quality of these estimates is dependent on the number of subjects that stay or switch and having a balance of people staying on the same treatment and switching treatments allows for better estimation of these parameters. The increased variance for the median cut off is also related to the correlation and $\alpha_m, m = 1$ and 3 parameter estimates. When these parameters and the correlation are fixed and do not need to be estimated, the bias is negligible and the variance in the bias estimates are consistent (Figure 2.5).

Figure 2.5: Average bias across all $\beta_j, j = A, B, C$ parameters for various dichotomized outcome (DO) cut offs when $\alpha_i = 0, i = 1, 3$ and $\tau_j = 0, j = 1, 2$. β_j parameters indicate treatment effect, $\alpha_i, i = 1, 3$ is the relationship between the treatments between stages, and τ is the correlation between outcomes for each stage. If a patient had a first stage outcome higher than the cut off, then the patient stayed on the treatment. Otherwise, they were re-randomized to one of the remaining treatments.

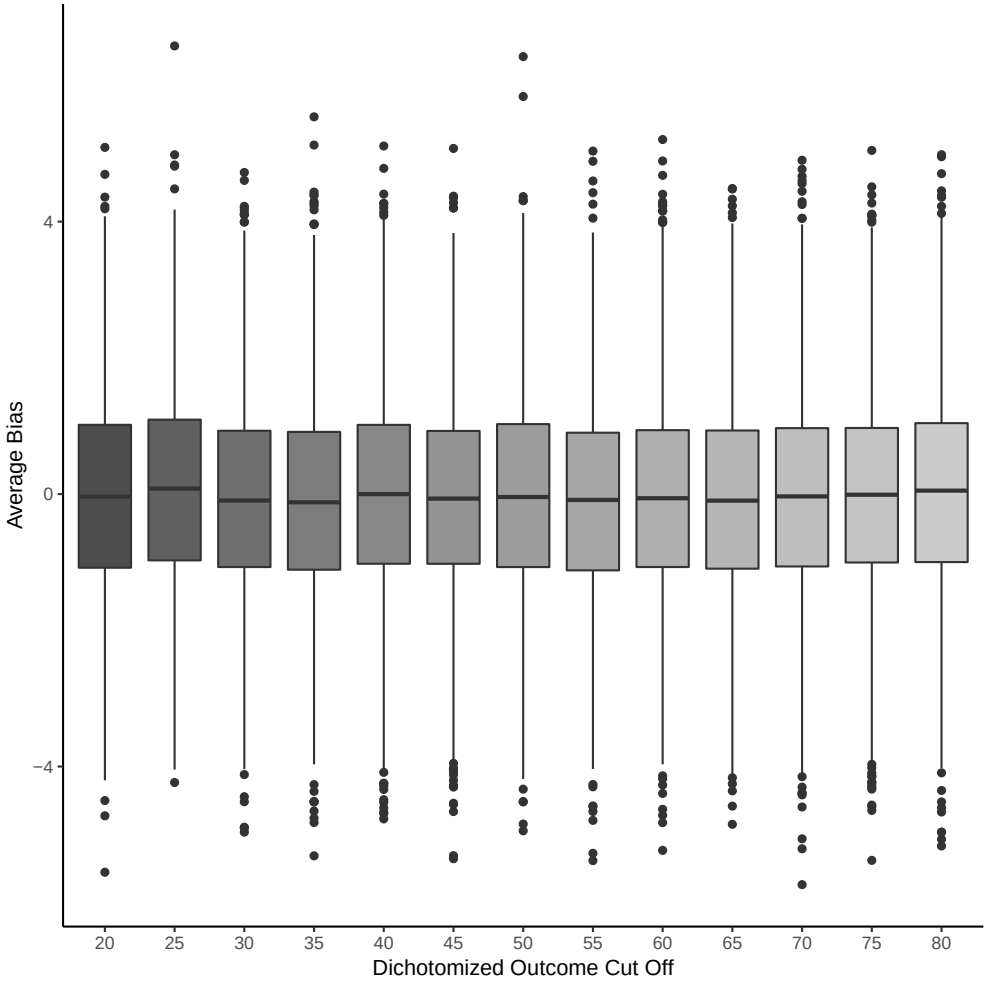

In general, the statistical performance was worse when using a binary outcome based on di-

chotomizing the outcome at the end of stage one than when using a continuous outcome with a mapping function. For some scenarios, binary outcomes performed similar to the mapping functions (Scenario 2, dichotomized outcome (DO) 70; Scenario 3, DO 30). However, the selection of the dichotomized outcome cut off influences the model parameter estimation which is not ideal since the optimal cut off value is not generally known in advance. For all scenarios, there were binary outcomes that performed worse than all mapping functions in terms of bias and efficiency. This indicates that poorly selected binary outcomes perform worse than any recommended mapping function. We additionally looked at mapping functions with a smaller range (25 to 75 instead of 0 to 100) and found that these functions also had improved performance over the dichotomized outcomes (Table 2.5).

Table 2.5: The root mean squared error (rMSE) for each $\beta_j, j = A, B, C$ when estimated in a single-stage design and the proposed two-stage snSMART design. β_j indicates the treatment effect. Here we use mapping functions (MF) where $Y_{min} = 25$ and $Y_{max} = 75$ but are other wise the same as those described in Section 2.3.1.

		rMSE		
Scenario	Design	β_A	β_B	β_C
1	MF 1	3.700	3.520	3.520
	MF 1/2	3.950	3.520	3.210
	MF 2	3.410	3.930	4.300
2	MF 1	2.990	3.320	3.750
	MF 2	2.820	3.020	3.440
	MF 1/2	3.220	3.470	3.780
3	MF 1	3.890	3.400	3.030
	MF 2	4.650	4.090	3.380
	MF 1/2	3.400	3.070	2.900

The percent of trials that correctly identified the optimal treatment were similar between mapping functions and dichotomous outcomes where the mapping function performed slightly better (Table 2.2, Table 2.6). In a null scenario where all treatments have treatment effect of 50, each treatment was identified as the optimal approximately 33% of the time which is as expected (Table 2.7).

One difference between snSMARTs with a mapping function and snSMARTs with a binary outcome was the proportion of subjects that stayed on the same treatment (Table 2.8). The proportion that stayed on the same treatment had less variation under a mapping function than that from a binary outcome. In scenarios where the dichotomization strategy was a poor fit for the patient outcomes (i.e. the cut off is either too high or too low), extreme results occurred where 3.2% stayed on treatment (Scenario 2, DO 70) or where 96.8% switched treatments (Scenario 3, DO 30). These same scenarios had somewhat less extreme results under a mapping function (MF)

Table 2.6: The proportion of trials where best treatment has the highest estimated first stage treatment effect when estimated under various scenarios. MF = mapping function, DO = dichotomized outcome, n = number of patients in one arm of the 3-arm trial.

		n		
Scenario	Design	15	30	100
1	MF 1	0.955	0.995	1
	MF 1/2	0.965	0.994	1
	MF 2	0.945	0.992	1
	DO 50	0.953	0.992	1
	DO 30	0.962	0.994	1
	DO 70	0.933	0.992	1
2	MF 1	0.947	0.993	1
	MF 1/2	0.963	0.996	1
	MF 2	0.928	0.991	1
	DO 50	0.924	0.987	1
	DO 30	0.959	0.994	1
	DO 70	0.913	0.969	1
3	MF 1	0.953	0.997	1
	MF 1/2	0.959	0.995	1
	MF 2	0.952	0.994	1
	DO 50	0.956	0.997	1
	DO 30	0.953	0.993	1
	DO 70	0.957	0.996	1

Table 2.7: The root mean squared error (rMSE) for each $\beta_j, j = A, B, C$ when estimated in a single-stage design and the proposed two-stage snSMART design. β_j indicates the treatment effect. Proportion chosen is the proportion of trials that estimated treatment j to have the highest treatment effect. MF = mapping function, DO = dichotomized outcome (Section 2.3.1).

		rMSE			Percent Chosen		
Scenario	Design	β_A	β_B	β_C	A	B	C
Null	MF 1	3	2.990	2.940	0.331	0.341	0.328
	MF $1/2$	2.890	2.900	2.980	0.345	0.331	0.324
	MF 2	3.160	3.110	3.150	0.324	0.336	0.340
	DO 50	6.540	6.550	6.550	0.316	0.335	0.349
	DO 30	4.120	4.140	4.140	0.325	0.327	0.348
	DO 70	3.610	3.640	3.670	0.322	0.343	0.335

where 13.5% and 85% stayed on treatment respectively (Scenario 2, MF 2; Scenario 3, MF 1/2). Since some model parameters may be impossible to estimate if a trial concludes with no patients on a certain treatment path (i.e. if all patients stay on the same treatment or all switch treatments) having a moderate proportion of individuals that stay on the same treatment is preferable to avoid such extreme circumstances. Additionally, patients all stay or switch is beneficial to the patients if all are doing very well or very poorly respectively. However, having all patients stay may prevent some patients from trying other treatments that perform better and having all switch may mean that the patients do not experience the long term benefits of staying on a treatment.

Table 2.8: The proportion of patients that stayed on the same treatment for ideal Scenarios 1-3 (Section 2.3.2.1) when using a dichotomized outcome (DO) compared to a mapping function (MF), $n = 30$.

	MF			DO		
Scenario	1	1/2	2	50	30	70
1	0.500	0.687	0.296	0.501	0.821	0.178
2	0.308	0.508	0.135	0.178	0.501	0.032
3	0.692	0.821	0.522	0.823	0.968	0.500

2.4.1.3 Model assumption violations

Here we examine the effects on the results when model assumptions are violated. Unlike the ideal scenarios, the bias heavily depended on the β parameter being estimated (Table 2.9, Figure 2.4B). We examine when the second stage treatment effect is based on the TSP of the first and second stage treatments rather than a weighted average (Scenarios 4, 7, 8, 10), variance is based on treatment and stage (Scenarios 5, 7, 9, 10), and correlation is based on the two treatments received (Scenarios 6, 8, 9, 10).

In general, the mean model violation caused the largest increase in bias (Figure 2.4B). Correlation and variance assumption violations did not result in increased bias suggesting that the estimates are some what robust to correlation and variance misspecification. Similarly, the coverage probability for scenarios with the mean model assumption violation was frequently less than the expected 0.95 (Table 2.10). The variance violation changed the variance of the estimates with estimates of β_A having smaller variance and estimates of β_C having larger variances (Figure 2.4B). Correlation assumption violations did not have a large impact on any statistical performance metric when not compounded with other assumption violations. Results were similar with other mapping functions (Table 2.13, Table 2.11). We also examined scenarios where the first stage outcome is a beta distribution with the same treatment effects as Scenarios 1, 2, and 3. These results show that

Table 2.9: The bias for each β when estimated under various scenarios. MF = mapping function, $\beta_j, j = A, B, C$ parameters indicate treatment effects.

		Bias		
Scenario	MF	β_A	β_B	β_C
1	MF 1	-0.248	-0.304	-0.283
	MF 1/2	0.247	0.124	0.089
	MF 2	-0.650	-0.791	-1.075
2	MF 1	-0.260	-0.427	-0.388
	MF 1/2	0.576	0.513	0.267
	MF 2	-0.278	-0.463	-0.758
3	MF 1	-0.166	-0.189	-0.130
	MF 1/2	-0.128	0.004	-0.191
	MF 2	-0.485	-0.688	-0.730

the mapping function has better statistical efficiency than the dichotomized outcome even when the model distribution is incorrectly specified. (Table 2.12).

Table 2.10: Performance results for Scenarios 4-10 with model assumption violations (Section 2.3.2.3), $n = 30$, linear mapping function (MF 1). $\beta_j, j = A, B, C$ parameters indicate treatment effects.

	Bias			rMSE		
Scenario	β_A	β_B	β_C	β_A	β_B	β_C
4	3.227	-1.117	-0.619	4.816	3.220	2.960
5	-0.247	-0.743	-1.261	1.741	3.206	4.722
6	-0.258	-0.276	-0.273	2.990	2.952	2.853
7	2.188	-2.158	-1.737	3.083	3.737	4.771
8	3.181	-1.146	-0.636	4.692	3.227	3.021
9	-0.257	-0.829	-1.228	1.777	3.212	4.604
10	2.214	-2.234	-1.824	3.101	3.890	4.775
	CP			CW		
Scenario	β_A	β_B	β_C	β_A	β_B	β_C
4	0.792	0.930	0.954	12.500	11.728	11.740
5	0.998	0.941	0.855	11.866	12.283	13.695
6	0.943	0.941	0.951	11.556	11.330	11.208
7	0.953	0.901	0.851	12.092	12.210	13.793
8	0.801	0.930	0.942	12.469	11.689	11.704
9	0.998	0.937	0.865	11.819	12.256	13.731
10	0.953	0.889	0.842	12.077	12.202	13.771

Table 2.11: Coverage probability (CP) and credible interval width (CW) for Scenarios 4-10 with model assumption violations (Section 2.3.2.3) for other mapping functions. MF = mapping function, $\beta_j, j = A, B, C$ parameters indicate treatment effects.

			CP			CW		
N	Scenario	MF	β_A	β_B	β_C	β_A	β_B	β_C
100	4	MF 1/2	0.265	0.922	0.938	6.873	6.484	6.367
		MF 2	0.867	0.906	0.935	6.738	6.510	6.775
	5	MF 1/2	0.999	0.943	0.851	6.769	6.869	7.270
		MF 2	0.999	0.830	0.424	6.346	6.486	7.430
	6	MF 1/2	0.941	0.937	0.935	6.518	6.341	6.266
		MF 2	0.922	0.910	0.901	6.219	6.179	6.354
	7	MF 1/2	0.414	0.870	0.872	6.870	6.894	7.304
		MF 2	0.975	0.790	0.788	6.539	6.681	7.981
	8	MF 1/2	0.258	0.915	0.940	6.859	6.473	6.354
		MF 2	0.859	0.913	0.942	6.735	6.514	6.763
	9	MF 1/2	1	0.955	0.856	6.779	6.861	7.259
		MF 2	0.999	0.808	0.425	6.349	6.491	7.439
	10	MF 1/2	0.445	0.878	0.885	6.864	6.908	7.303
		MF 2	0.969	0.796	0.774	6.546	6.655	7.972
30	4	MF 1/2	0.692	0.928	0.925	12.182	11.504	11.340
		MF 2	0.888	0.920	0.944	12.201	11.604	12.156
	5	MF 1/2	0.999	0.958	0.875	11.877	12.121	12.882
		MF 2	0.998	0.914	0.725	11.397	11.738	13.398
	6	MF 1/2	0.942	0.936	0.944	11.514	11.217	11.035
		MF 2	0.930	0.933	0.928	11.277	11.191	11.269
	7	MF 1/2	0.899	0.904	0.856	12.040	12.157	12.908
		MF 2	0.988	0.878	0.809	11.816	11.858	14.253
	8	MF 1/2	0.701	0.918	0.939	12.184	11.493	11.323
		MF 2	0.894	0.929	0.947	12.256	11.657	12.199
	9	MF 1/2	0.998	0.958	0.876	11.802	12.095	12.810
		MF 2	0.998	0.918	0.718	11.436	11.752	13.425
	10	MF 1/2	0.896	0.893	0.847	12.042	12.061	12.842
		MF 2	0.983	0.899	0.804	11.853	11.898	14.278
15	4	MF 1/2	0.827	0.925	0.916	16.716	15.853	15.653
		MF 2	0.885	0.920	0.935	16.991	16.110	16.716
	5	MF 1/2	0.994	0.935	0.848	15.890	16.518	17.818
		MF 2	0.998	0.920	0.789	15.708	16.217	18.582
	6	MF 1/2	0.933	0.932	0.930	15.889	15.541	15.274
		MF 2	0.926	0.921	0.921	15.662	15.489	15.497
	7	MF 1/2	0.953	0.889	0.827	16.208	16.572	17.929
		MF 2	0.974	0.897	0.827	16.159	16.237	19.274
	8	MF 1/2	0.832	0.919	0.923	16.685	15.858	15.615
		MF 2	0.883	0.926	0.945	16.965	16.077	16.669
	9	MF 1/2	0.995	0.934	0.843	15.893	16.508	17.775
		MF 2	0.997	0.924	0.772	15.723	16.227	18.681
	10	MF 1/2	0.947	0.889	0.836	16.195	16.525	17.981
		MF 2	0.973	0.903	0.825	16.190	16.292	19.377

Table 2.12: The root mean squared error (rMSE) for each $\beta_j, j = A, B, C$ when estimated in the proposed two-stage snSMART design when the first stage outcome is distributed as a scaled beta distribution, but analyzed as a normal distribution. β_j indicates the treatment effect. MF = mapping function, DO = dichotomized outcome (Section 2.3.1).

		rMSE		
Scenario	Design	β_A	β_B	β_C
Beta - 1	MF 1	4.610	4.440	4.420
	MF 1/2	4.360	4.220	4
	MF 2	4.250	4.430	4.470
	DO 50	10.100	9.630	9.300
	DO 30	8.710	7.860	7.200
	DO 70	5.660	6.330	6.860
Beta - 2	MF 1	6.100	6.800	6.880
	MF 2	4.890	5.260	5.650
	MF 1/2	7.530	8.140	8.380
	DO 50	8.950	10.780	12.720
	DO 30	13.220	17.980	22.020
	DO 70	5.490	5.940	6.800
Beta - 3	MF 1	10.070	9.140	12.790
	MF 2	16.130	13.770	14.900
	MF 1/2	6.740	6.300	11.250
	DO 50	16.270	13.240	15.700
	DO 30	9.500	8.420	13.190
	DO 70	25.610	19.970	18.020

Table 2.13: Bias and root mean squared error (rMSE) for Scenarios 4-10 with model assumption violations (Section 2.3.2.3) for other mapping functions, $n = 30$. $\beta_j, j = A, B, C$ parameters indicate treatment effects.

		Bias			rMSE		
Scenario	MF	β_A	β_B	β_C	β_A	β_B	β_C
4	MF 1/2	4.331	−1.404	−1.295	5.417	3.192	3.145
	MF 2	1.587	−1.075	0.020	3.814	3.298	3.102
5	MF 1/2	−0.163	−0.034	0.259	1.753	2.953	4.189
	MF 2	−0.199	−1.656	−3.882	1.676	3.426	5.938
6	MF 1/2	0.043	0.199	0.078	2.996	3.016	2.901
	MF 2	−0.607	−0.880	−1.014	3.081	3.037	3.154
7	MF 1/2	3.270	−2.137	−1.356	3.890	3.677	4.289
	MF 2	1.282	−2.155	−2.953	2.469	3.721	5.430
8	MF 1/2	4.280	−1.398	−1.380	5.395	3.229	3.068
	MF 2	1.594	−1.138	0.011	3.801	3.195	3.127
9	MF 1/2	−0.108	−0.077	0.167	1.777	2.965	4.089
	MF 2	−0.242	−1.680	−3.964	1.723	3.417	6.068
10	MF 1/2	3.283	−2.127	−1.475	3.907	3.708	4.414
	MF 2	1.367	−1.988	−2.925	2.544	3.645	5.475

2.4.2 Patient outcomes

From the patient perspective when considering whether to enroll on a clinical trial, expected outcomes for the patients on that trial are a key consideration. Since the first stage of the trial with the mapping function and with the dichotomized outcome are the same, we focus on the second stage outcomes. Overall the outcomes between mapping functions and dichotomized outcomes were very similar within scenarios (Figure 2.6). Additionally, the patient outcomes did not vary by which mapping function was used indicating that the average patient experience in such a trial is not heavily influenced by the selection of the mapping function.

2.5 Discussion

We demonstrated that mapping functions are a flexible way to design multistage trials and this design can be used to efficiently estimate treatment effects while achieving similar patient outcomes. We see that the most intuitive choice for the binary outcome, dichotomizing a continuous outcome by selecting a cut off at the median treatment effect, has low bias but also has high variance. Other cut offs had high bias, but lower variance. The mapping function does not have this extreme bias-variance trade off and can provide efficient and unbiased estimates in small samples. Although sample size calculations are outside of the scope of this chapter, we ex-

Figure 2.6: Median second stage outcome across all simulated trials. Lines indicate the 25th and 75th percentile of second stage outcomes across all simulated trials. MF = mapping function (Section 2.3.1), DO = dichotomized outcome (Section 2.3.2.2)

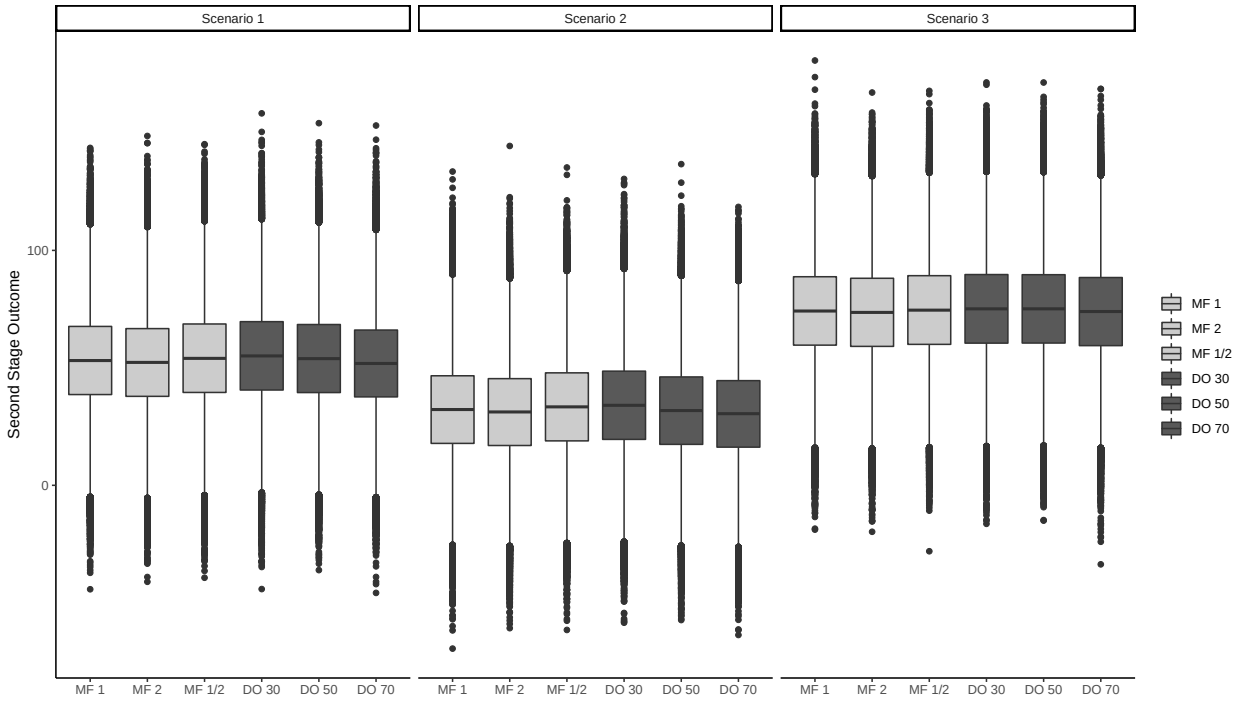

pect that the mapping function snSMART design will typically require smaller sample size than a binary outcome snSMART due to the smaller rMSE. Our group has developed sample size calculations for an snSMART with binary endpoint (Wei et al., 2020) and a corresponding applet (https://umich-biostatistics.shinyapps.io/snsmart_sample_size_app). If a more refined sample size estimate is needed, simulation is always possible and the R code that generated the simulation results in this chapter is available from the corresponding author.

Mapping functions could take many forms. While we primarily discussed linear functions and modifications to those functions, there may be cases where more is known about the distribution of Y_1 which could be used to generate a mapping function. For instance, Robson et al. (2018) describes the distribution of a PRO for ANCA-associated vasculitis for both when the patient has active disease or is in remission. In this case, the empirical distribution of this PRO for patients in remission could be used as a mapping function. A similar approach could be used for when the PRO or other continuous outcome is well studied and has a known empirical distribution (Wei et al., 2000). Using the empirical cumulative distribution results in the median value of Y_1 having equal probability of staying or switching treatments. In all distributions, even when heavily skewed, this mapping function has the trait such that if a patient has an outcome above the median, they have a higher probability of staying on treatment and if the outcome is below the median they have a higher probability of switching treatment, but is more efficient in terms of rMSE than using a median cut off.

While we focused on a continuous outcome in a small sample setting, this trial design is very flexible in terms of the number of treatments, Bayesian analysis, and distribution of the outcome. We presented the design with three treatments, but it is easily extendable to more treatments. The probability of staying on the same treatment would still be determined by the mapping function. If a patients switches treatments, the second treatment would be randomly selected from the remaining treatments with each remaining treatment having equal probability. The Bayesian analysis framework is also very adaptable. Baseline measures and covariates measured up until the re-randomization can be added to the mean models for stages one and two. Other covariates measured during the second stage of the trial up until the final outcome measurement can be added to mean model in equation 2.1. However, with small samples, parsimony in the model is important. Likewise, if treatment specific pathway effects are the primary interest, the mean model could be changed to be unique for each treatment regimen rather than a parsimonious model using weighted means. Since an incorrectly specified mean model resulted in the largest increase in bias, we suggest sensitivity analysis of the mean model specification. Similarly, if the distributions of the outcome are not expected to be normal, other multivariate distributions can be used for the likelihood to be more suited to the outcome of interest.

It is not surprising that we did not see improved patient outcomes when using a mapping func-

tion relative to a binary outcome. This design is primarily focused on improving statistical efficiency while maintaining similar outcomes for patients enrolled in the trial. One possibility to improve patient outcomes would be to incorporate an adaptive re-randomization scheme where prior patient outcomes are included in the mapping function in addition to the patient's own outcome. This could be similar to the non-dichotomous randomized play-the-winner proposed in Rosenberger (1993).

This trial design still has many benefits to patients. Patient input could help guide the selection of the outcome or determine the minimum and maximum values where switching treatment and remaining on treatment are possible. Rare disease patient representatives stated that increasing patient involvement would improve clinical trials (Gaasterland et al., 2019). Additionally, patients have the potential to receive two treatments or stay on treatment if it is working. This lessens time on treatments that are less effective for the patient which may be viewed as a benefit to patients (Gaasterland et al., 2019). Informing patients about the clinical trial procedure and ensuring that patients understand that they may not receive the treatment they prefer due to randomization is also critical (Gaasterland et al., 2019). While the mapping function increases the probability of switching if a treatment is not working and staying if it is, there is no guarantee since there is randomization for all treatment assignments. To further patient involvement, the upper and lower values of re-randomization could potentially vary for each patient depending on their goals and needs. Future work will investigate adaptive randomization components to potentially improve patient outcomes and to consider individualized mapping functions to fully incorporate patient-specific willingness to stay on or switch treatment.

Overall, this snSMART design with a mapping function has improved statistical efficiency and lower bias over using a dichotomized outcome. The proposed design with a mapping function allows trials to proceed without the need for a pilot study or extensive knowledge of the treatment effects and outcome and may be particularly useful in estimating treatment effects in small samples.

CHAPTER 3

A Sequential, Multiple Assignment, Randomized Trial Design with a Tailoring Function

3.1 Introduction

Dynamic treatment regimens (DTRs) are a sequence of decision rules that provide guidelines for treatments based a patient's their individual characteristics and disease progression (Chakraborty and Murphy, 2014). These treatment sequences are a form of personalized medicine where patients with the same condition may receive different treatments based on the evolution of the condition, additional patient characteristics, and response to prior treatments. A decision rule tailors treatment at each time point for each patient by implementing a tailoring variable (Almirall et al., 2012; Lei et al., 2012). An example of a two-stage DTR would be "Start all patients on treatment A. After 6 months, if the patient responds, maintain treatment A. Otherwise, give the patient treatment B." This DTR could be written as the treatment triplet (A, A, B) to indicate the first treatment and then the second two treatment options based on response to the initial treatment. The tailoring variable here is response to treatment A where response would have a predetermined definition.

One way to develop DTRs is through observational studies (Moodie et al., 2012; Murphy et al., 2001; Tao et al., 2018). In observational studies, often the second stage treatment may be chosen based on multiple factors and not a single binary variable as presented in the example above (i.e. repsonse/non-response to treatment). Thus, methods to estimate DTRs from observational data must determine the categorical variables that dictate which treatment is optimal at each time point.

Another way to develop DTRs is through randomized trials, most often sequential, multiple assignment, randomized trials (SMARTs). SMARTs integrate the tailoring variable into the design to assign a second stage treatment based on the patient's response to the previous treatment (Murphy, 2005). In SMARTs, DTRs are embedded into the design so that the effects of each DTR can be estimated from the resulting data. In a two-stage SMART, participants are randomly assigned to a first treatment and a first stage outcome or tailoring variable is measured. The second treatment is

assigned probabilistically or deterministically based on the design and the selected tailoring variable. The final outcome is measured at the end of the second stage. Data from both stages are used to estimate DTRs using weighted and replicated regression (WRR) with generalized estimating equations (GEE) (Kidwell et al., 2018; Murphy, 2005; Nahum-Shani et al., 2012a) or more personalized DTRs including baseline and time varying covariates can be developed using Q-learning (Nahum-Shani et al., 2012b).

The tailoring variable in a SMART is often based on response or non-response to treatment, thus requiring that the trial protocol specify a binary rule for what qualifies as a "response." The tailoring variable must be selected and agreed upon before the start of the trial (Lei et al., 2012). While there is some guidance in selecting the tailoring variable (Almirall et al., 2012), there is a lack of guidance on how to proceed if a reasonable tailoring variable cannot be established. In many studies, the outcome of interest may not be a binary variable and selecting a tailoring variable may not be an easy or practical task.

In this chapter, we propose a SMART design where the second stage treatment is assigned based on a tailoring function. A tailoring function, similar to a mapping function in Hartman et al. (2021) and Chapter 2, takes a continuous outcome at the end of the first stage and assigns a probability of receiving each possible second stage treatment. In the design presented here, we assume a binary response tailoring variable is not available, and thus traditional methods for analyzing data from a SMART are no longer applicable. Therefore, to develop DTRs, we apply methods for developing DTRs using observational data such as Q-learning and tree-based reinforcement learning (TBRL). This project addresses a gap in clinical trial methodology by presenting SMARTs where second stage treatment is based on a continuous outcome.

In this chapter, we first present the proposed SMART design, a SMART which uses a tailoring function (TF-SMART), with a goal of estimating a DTR in flexible manner without needing to pre-specify a tailoring variable. Next, we discuss several methods for developing DTRs from observational studies in the context of a SMART thereby eliminating the need for a pilot study to select a definition of response and non-response. Lastly, we show results from simulations to compare the estimated DTR from a TF-SMART to a traditional SMART using a predefined tailoring variable (TV-SMART). This new SMART design will simultaneously define decision rules and the appropriate sequence of treatments within the context of randomized trial, a novel approach to using randomized trials to develop DTRs.

3.2 Trial designs

We examine two TF-SMART designs presented in Figures 3.1 and 3.2. In these designs, patients are first randomized with equal probability to one of two treatments. Then, a continuous outcome is

measured at the end of the first stage, Y_1. The second stage treatment is assigned using a tailoring function (more details in Section 3.2.1). The tailoring function takes Y_1 as input and outputs a probability. The output probability is then used to assign probabilities of the patient receiving each second stage treatment. In Figures 3.1 and 3.2, we present distinct second stage treatments for generalizability, but the first stage treatments could be treatment options in the second stage. Second stage treatment options may include continuing the same treatment, alternate treatments, a change in dosage, a combination of treatments, or maintenance therapy. Examples of these designs with first stage treatments included as options for second stage treatments are shown in Figures 3.3 and 3.4.

Figure 3.1: Design I: Patients with more favorable responses are more likely to go to one treatment and patients with less favorable responses are more likely to go to a different single treatment. A, B, C, D, E, and F are treatments, Y_1 is the first stage outcome, R indicates a randomization point, $f(Y_1)$ is the tailoring function, and the expressions along the lines indicates the randomization probability of the treatment following.

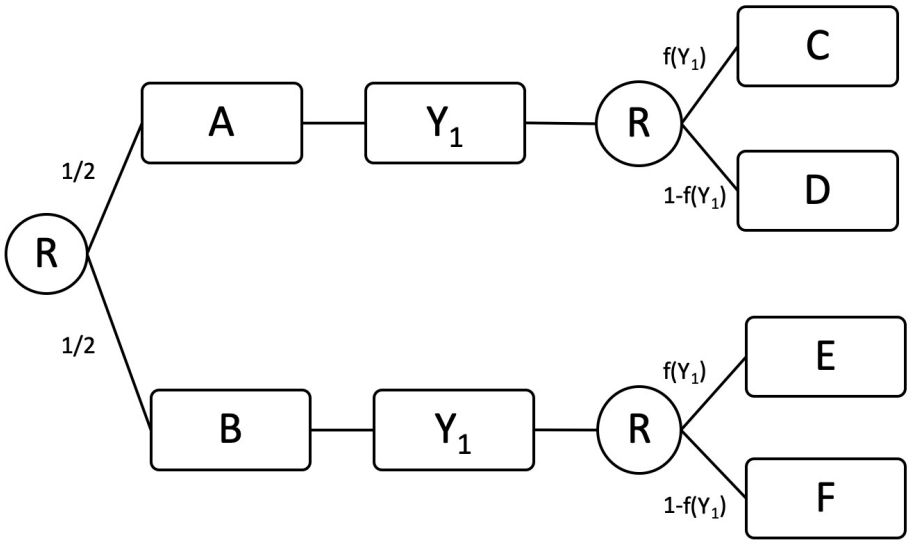

In Design I (Figure 3.1), there are two first stage treatments, A and B, and two second stage treatments for each arm, C and D or E and F. At the end of the first stage, a patient is randomly assigned to one of the two second stage treatments based on Y_1. When a tailoring variable is used in place of a tailoring function, then patients are assigned to either C/E or D/F based on their response to the first stage treatment and there is no second randomization point. In this case,

Figure 3.2: Design II: Patients with more favorable responses are more likely to go to one specific treatment and those with less favorable responses are more likely to be re-randomized to two treatments. A, B, C, D, E, F, G, and H are treatments, Y_1 is the first stage outcome, R indicates a randomization point, $f(Y_1)$ is the tailoring function, and the expressions along the lines indicates the randomization probability of the treatment following.

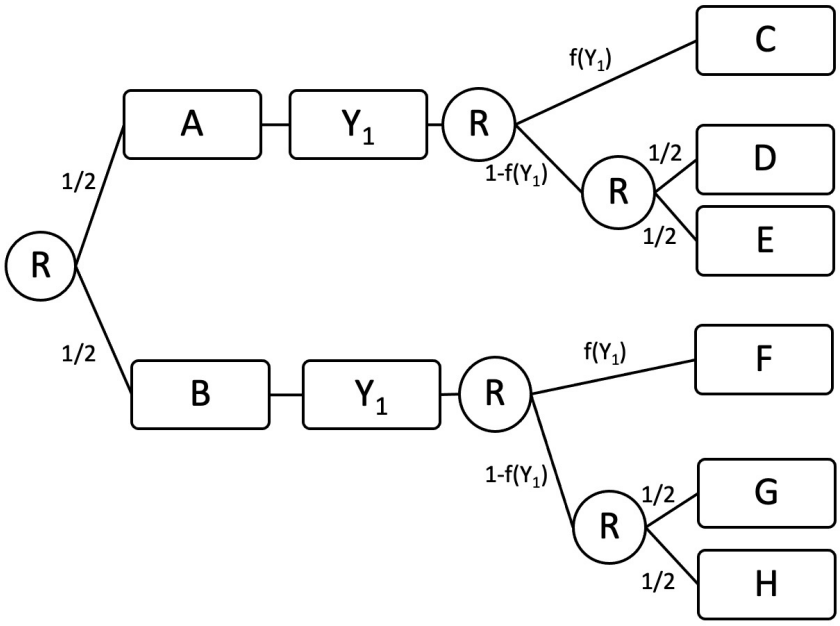

Figure 3.3: Example of Design I: Patients with more favorable responses are more likely to stay on the same treatment and patients with less favorable responses are more likely to add the treatment they did not receive in the first stage. A and B are treatments, Y_1 is the first stage outcome, R indicates a randomization point, $f(Y_1)$ is the tailoring function, and the expressions along the lines indicates the randomization probability of the treatment following.

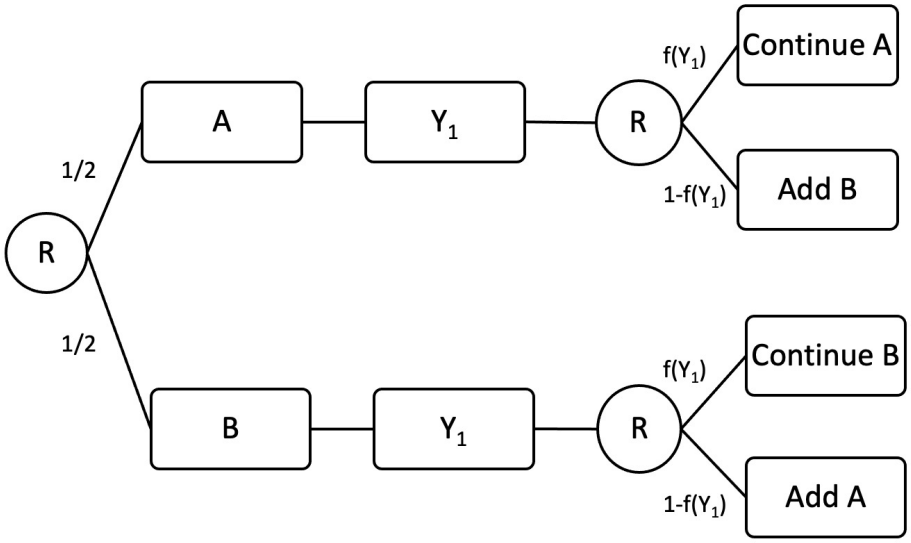

Figure 3.4: Example of Design II: Patients with more favorable responses are more likely to stay on the same treatment and patients with less favorable responses are more likely to either add the treatment they did not receive in the first stage or switch to a different treatment. A, B, and C are treatments, Y_1 is the first stage outcome, R indicates a randomization point, $f(Y_1)$ is the tailoring function, and the expressions along the lines indicates the randomization probability of the treatment following.

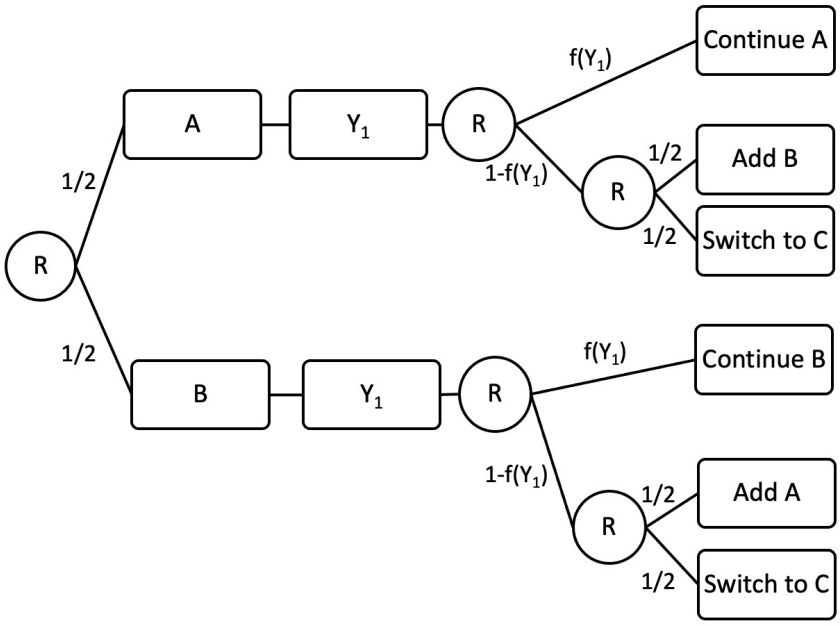

$f(Y_1) = 1$ or 0 so the second stage treatments are assigned deterministically based on Y_1. Thus, when a tailoring variable is used, this design is not a SMART as there is only one randomization point, not sequential randomization.

In Design II (Figure 3.2), patients are assigned a second stage treatment in two steps. First, patients are assigned to either receive one treatment, C or F in the figure, or receive one of the two other treatments, D and E or G and H with equal probability. One common example of this general design occurs when C = A and F = B so these second stage treatments may be described as continuing first stage treatment or some form of maintenance care. The tailoring function would be defined such that a patient who is doing well on the first stage treatment would have a higher probability of continuing that treatment or receiving a maintenance treatment.

3.2.1 Tailoring function details

Tailoring functions are a generalization of tailoring variables and are mechanistically the same as mapping functions presented in Hartman et al. (2021) and Chapter 2. However, interest in this chapter is in developing DTRs in a standard size trial which differs from previous work where the goal is to estimate the best first stage treatment in a small sample. A tailoring function takes a continuous first stage outcome, Y_1, and maps it to randomization probabilities of each second stage treatment. For designs such as those in Figures 3.3 and 3.4, this function can map better outcomes to values closer to 1 so that patients are more likely to stay on the treatment that worked well for them. Conversely, patients that are not doing well on the first stage treatment are more likely to switch to a new treatment or add an additional treatment with the goal of improving the second stage outcome. There are many options for selecting a function. Guidance from Section 2.2.1.1 can also apply here. As in Chapter 2, the function itself can implement thresholds to limit extreme circumstances, and the function can be linear or otherwise to modify its properties.

Tailoring variables can be thought of as a special case of tailoring functions where $f(Y_1) = 0$ for $Y_1 < x$ and $f(Y_1) = 1$ for $Y_1 \geq x$ for some x which is the pre-selected cut off. Using a tailoring variable defined as a step-wise function in this manner generates a TV-SMART design with a binary first stage outcome that assigns second stage treatment.

3.3 Overview of DTR estimation methods

The goal of a DTR is to maximize the expected second stage outcome when applied to the population. In a traditional SMART design with a binary intermediate outcome, WRR with GEE is commonly used to analyze data and estimate a DTR (Kidwell et al., 2018; Liang and Zeger, 1986; Murphy, 2005; Nahum-Shani et al., 2012b). This requires a predefined tailoring variable as there

is no way to develop a tailoring variable within the WRR framework. Q-learning may also be used to analyze data from a SMART to develop more personalized DTRs which include baseline and time varying covariates (Nahum-Shani et al., 2012b).

Published methods for DTR estimation from observational data include outcome weighted learning (OWL) (Zhao et al., 2012), Q-learning (Moodie et al., 2012; Nahum-Shani et al., 2012b; Wu et al., 2015), TBRL (Laber and Zhao, 2015; Tao et al., 2018), and Bayesian machine learning (Murray et al., 2018). Here, we present WRR in the context of a TV-SMART design and we present Q-learning and TBRL in the context of the TF-SMART designs shown in Figures 3.1 and 3.2.

Let Y_1 and Y_2 be the first and second stage outcomes, respectively. Let A_1 and A_2 be the first and second stage treatment assignments, respectively. The goal of a SMART is to identify the DTR which maximizes the expected second stage outcome, Y_2, for patients in a new population.

3.3.1 Generalized estimating equations

The WRR method to analyze SMART data weights observations based on their probability of receiving the treatment regimen they received. In a traditional SMART, there are multiple DTRs embedded within the trial design. In some cases, people may be on a sequence of treatments which could have resulted due to following multiple DTRs. Thus, those patients must be replicated and each replicate is assigned one of the possible DTRs they could be following. Patients are also weighted to account for the randomization scheme. Then a fully saturated model is often used. The weighting and replicating method allows the data to be analyzed by standard software (Kidwell et al., 2018; Nahum-Shani et al., 2012a). Due to our proposed TF-SMART design not having a pre-specified decision rule for defining a binary tailoring variable, this method cannot be used to estimate DTRs for our proposed TF-SMART design as WRR only estimates DTRs with a predefined tailoring variable. The analytic methods we use must both develop the tailoring variable and define the treatment regimen, whereas WRR assumes a predetermined tailoring variable and only defines the treatment regimen. We use this analytic method to analyze data from a TV-SMART and compare results with a TF-SMART using other methods.

3.3.2 Q-learning

Q-learning is a method to maximize an outcome based on a sequence of decisions (Nahum-Shani et al., 2012b). First, the optimal decision rule at the second stage is determined by identifying A_2 that maximizes $Q_2(A_1, A_2, Y_1) = E(Y_2|A_1, A_2, Y_1)$. Then, the optimal first stage decision rule is constructed by selecting A_1 that maximizes $Q_1(A_1) = E\left(\max_{A_2} Q_2(A_1, A_2, Y_1)\right)$. The quantity $E\left(\max_{A_2} Q_2(A_1, A_2, Y_1)\right)$ is called a pseudo-outcome. This represents the ex-

pected second stage outcome if the patient received the best second stage treatment. Q-learning can include covariates measured before the first stage (X_1) and up to the second stage (X_2) so that $Q_2(A_1, A_2, Y_1, X_1, X_2) = E(Y_2|A_1, A_2, Y_1, X_1, X_2)$ and $Q_1(A_1, X_1) = E\left(\max_{A_2} Q_2(A_1, A_2, X_2, X_2, Y_1)\right)$. Q_1 and Q_2 are called Q-functions. While Q-learning can be done with a variety of techniques and models, here we focus on linear regression so that Q_2 is modeled as a linear combination of A_1, A_2, and Y_1 and Q_1 is modeled as a linear combination of A_1.

In order to account for the randomization scheme, we use inverse probability of treatment weighting (IPTW) (Moodie et al., 2012). For each patient, the probability of their treatment is the product of the probability of first stage treatment assignment ($\frac{1}{K}$ for K first stage treatments assuming equal randomization) and the probability of their second stage treatment which is dictated by the tailoring function for that patient and the randomization scheme.

3.3.3 Tree-based reinforcement learning

We follow methods from Tao et al. (2018) and Laber and Zhao (2015) for TBRL using IPTW in the TF-SMART setting. As with Q-learning, we start by identifying the optimal second stage treatment and then use a pseudo-outcome to identify the optimal first stage treatment. With TBRL, a tree is built to create decision rules for the second stage treatment assignment that maximizes the second stage outcome. Then, a tree is built for the first stage treatment assignment to maximize the expected second stage outcome using pseudo-outcomes.

Decision trees start with all patients and use binary rules to create groups of patients that should all receive the same treatment. An example is shown in Figure 3.5. Here, the blue and green squares are termed nodes. The blue squares are terminal nodes. Nodes are terminal when no additional split follows. The orange squares are the treatments assigned to each terminal node.

To build a tree, we first define a node purity measure which represents how well a tree has classified the groups and is a measure of homogeneity within a node. We use the purity measure described in (Laber and Zhao, 2015):

$$\mathcal{P}(r, \mathcal{R}) = \max_{a,a' \in \mathcal{A}} E_n \left[\frac{[Y - \hat{m}(X)]\, I(X \in \mathcal{R}) I(A = \pi_{ra,a'}(X))}{P(\pi_{r,a,a'}(X)|X)} \right] \left(E_n \left[\frac{I(X \in \mathcal{R}) I(A = \pi_{r,a,a'}(X))}{P(\pi_{r,a,a'}(X)|X)} \right] \right)^{-1} \quad (3.1)$$

where A is the treatment, $\mathcal{R}$ is the node to be split into groups r and r^C, X are patient characteristics for that specific stage, $\pi_{r,a,a'}$ is the decision rule that assigns treatment a to subjects

Figure 3.5: An example of a potential decision tree for a cancer patient. The blue and green squares are nodes. The blue squares are terminal nodes. The orange squares are the treatments estimated to be optimal for patients in that node.

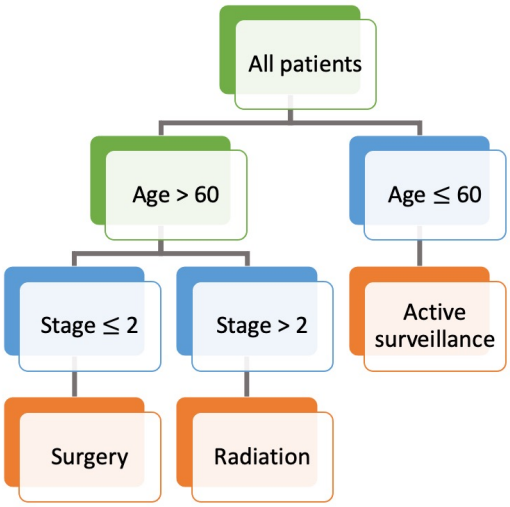

in group r and treatment a' to subjects in r^C, and $\hat{m}(X)$ is $\max_{a \in \mathcal{A}} E(Y|A = a, X)$. E_n is the empirical expectation. Then $P(\pi_{r,a,a'}(X)|X)$ is the propensity score for the patient defined by the inverse of the probability of receiving the treatment that they received. This purity measure is the percent change in the expected outcome when the patients in a given node are split (i.e. patients in r receive treatment a and those in r^C receive a') relative to the expected outcome if patients in that node are not split (i.e. all patients in the node receive the same treatment).

One distinct difference in this TF-SMART setting versus the observational data setting is that we do not need to estimate propensity score since the probability of each treatment is a known function of the first stage outcome via the tailoring function and randomization probability.

With the node purity measure, we then define splitting rules. Each split must only depend on one covariate. In the scenario without covariates, this means that there will be at most one split which is solely based on the first stage outcome. If there are other patient covariates that could be considered, then the split is based on one out of all covariates measured up until randomization. The split is chosen to optimize the purity metric defined in equation 3.1. Node size is defined as the number of patients in that node. Define n_0 as the minimum node size accepted where n_0 is a positive integer. Define λ as the minimum improved node purity to warrant a split where λ is a positive constant. Then we define our splitting rules as follows:

Rule 1 If a node size is smaller than $2n_0$, the node will not be split.

Rule 2 If all possible splits result in a child node having fewer than n_0 items, the node will not be

split.

Rule 3 If the tree-depth (or number of splits) has reached a pre-specified limit, the node will not be split.

If these criteria are met, then the node is split by selecting r such that $\mathcal{P}(r, \mathcal{R})$ is maximized and the node size of r and r^C are both at least size n_0. Let $|\cdot|$ indicate node size. Then the optimal split is:

$$\hat{r}^{opt} = \arg\max_r [\mathcal{P}(r, \mathcal{R}) : \min(|r|, |r^C|) \geq n_0]$$

Once this optimal split has been determined, we apply the last rule:

Rule 4 If the purity improvement is less than λ (i.e. $\mathcal{P}(r, \mathcal{R}) - \mathcal{P}(\mathcal{R}) < \lambda$), then do not split the node.

This splitting rule forces the expected outcome to improve by some set amount in order to warrant a split. We follow these splitting rules until all nodes are terminal (i.e. no split follows all rules defined above) which results in an estimate of the optimal DTR.

3.4 Simulation Methods

3.4.1 Generation of trial data

Patients are randomly assigned to one first stage treatment so that there are equal numbers of patients in all treatment arms. Assume there are s possible first stage treatments and t possible second stage treatments. Then all possible first stage treatments are $\mathbf{A_1} = \{a_{11}, \ldots, a_{1s}\}$ and all possible second stage treatments are $\mathbf{A_2} = \{a_{21}, \ldots, a_{2t}\}$. Then the first and second treatments a patient receives (A_1 and A_2 respectively) are such that $A_1 \in \mathbf{A_1}$ and $A_2 \in \mathbf{A_2}$. The first stage outcome is distributed:

$$Y_1|A_1 \sim N\left(\sum_{i=1}^{s} \alpha_i I(A_1 = a_{1i}), \sigma^2\right)$$

Each treatment $a_{11}, \ldots, a_{1s}$ has its own mean treatment effect on the first stage outcome. Here, since $s = 2$ such that $a_{11} = A$ and $a_{12} = B$, the distribution of Y_1 is:

$$Y_1|A_1 \sim N\left(\alpha_A I(A_1 = A) + \alpha_B I(A_1 = B), \sigma^2\right).$$

The second stage treatment is randomly assigned using a tailoring function based on the first stage outcome, Y_1 and the first stage treatment, A_1. Let $\boldsymbol{\pi}(Y_1, A_1)$ be the vector of probabilities

of each second stage treatment as dictated by the trial design and the tailoring function. Then the second stage treatment assignment is:

$$A_2|Y_1, A_1 \sim Multinomial(\boldsymbol{\pi}(Y_1, A_1)).$$

The distribution of the final outcome Y_2, is dependent on all previous treatments and the first stage outcome. We assume a linear relationship between Y_1 and Y_2 given A_1 and A_2. Thus the distribution of Y_2 is:

$$Y_2|A_2, Y_1, A_1 \sim Normal(\mu(A_2, Y_1, A_1), \sigma^2)$$

where:

$$\mu(A_2, Y_1, A_1) = \sum_{i=1}^{s}\sum_{j=1}^{t} \beta_{a_{1i}a_{2j}} I(A_1 = a_{1i})I(A_2 = a_{2i}) + \sum_{i=1}^{s}\sum_{j=1}^{t} \gamma_{a_{1i}a_{2j}} I(A_1 = a_{1i})I(A_2 = a_{2i})Y_1. \quad (3.2)$$

This model is a general formulation and there are more parsimonious models that would be special cases of this model by imposing additional constraints on the γ and/or β parameters.

3.4.2 Analysis methods

3.4.2.1 Weights

For both Q-learning and TBRL, we use IPTW. Weights for the first stage analyses in both TF-SMART Designs I and II (Figures 3.1 and 3.2) are $w_1 = \frac{1}{2}$ due to the equal stage randomization to the two first stage treatments. Since all patients are weighted equally, we do not need to weight at this stage. The weights for patients vary in the second stage and are defined for Design I as:

$$w_2 = \frac{2}{f(y_1)} I(A_2 \in (C, E)) + \frac{2}{(1 - f(y_1))} I(A_2 \in (D, F)) \quad (3.3)$$

and Design II as:

$$w_2 = \frac{2}{f(y_1)} I(A_2 \in (C, F)) + \frac{4}{(1 - f(y_1))} I(A_2 \in (D, E, F, G)). \quad (3.4)$$

3.4.2.2 Q-Learning

To implement Q-learning in our design settings, we use linear models to identify the optimal treatment at each randomization point in the specific design. For Design I, this means we separately identify which is the best second stage treatment for patients that received treatment A in the first stage (C or D) and the best second stage treatment for patients that received treatment B in the first stage (E or F). In Design II, we identify the best second stage treatments for those that received A in the first stage (C, D, or E) and those that received B in the first stage (F, G, or H). Lastly, in both designs, we identify which is the best treatment to start with to maximize the second stage outcome (A or B). Let β_{ij} and γ_{ij} be the model parameters for when $A_1 = i$ and $A_2 = j$. The model to identify the best second stage treatment for those starting with treatment A in Design I is:

$$Y_2 \sim \beta_{AC} I(A_2 = C) + \beta_{AD} I(A_2 = D) + \gamma_{AC} Y_1 I(A_2 = C) + \gamma_{AD} Y_1 I(A_2 = D) \tag{3.5}$$

Likewise, the model for identifying the best second stage treatment for those starting with treatment B is:

$$Y_2 \sim \beta_{BE} I(A_2 = E) + \beta_{BF} I(A_2 = F) + \gamma_{BE} Y_1 I(A_2 = E) + \gamma_{BF} Y_1 I(A_2 = F) \tag{3.6}$$

The equivalent models for Design II are:

$$\begin{aligned} Y_2 \sim \beta_{AC} I(A_2 = C) + \beta_{AD} I(A_2 = D) + \beta_{AE} I(A_2 = E) + \\ \gamma_{AC} Y_1 I(A_2 = C) + \gamma_{AD} Y_1 I(A_2 = D) + \gamma_{AE} Y_1 I(A_2 = E) \end{aligned} \tag{3.7}$$

and

$$\begin{aligned} Y_2 \sim \beta_{BF} I(A_2 = F) + \beta_{BG} I(A_2 = G) + \beta_{BH} I(A_2 = H) + \\ \gamma_{BF} Y_1 I(A_2 = F) + \gamma_{BG} Y_1 I(A_2 = G) + \gamma_{BH} Y_1 I(A_2 = H). \end{aligned} \tag{3.8}$$

We use the linear regression with the above models and IPTW with the appropriate weights from equation 3.3 or 3.4 to obtain estimates for all βs and γs.

Let $\hat{Y}_2$ indicate the predicted second stage outcome based on the parameter estimates from the models above. Then the pseudo outcome to analyze the stage 1 data for Design I is:

$$\tilde{Y}_2 = \max_{A_2 \in C,D} \hat{Y}_2 I(A_1 = A) + \max_{A_2 \in E,F} \hat{Y}_2 I(A_1 = B) \tag{3.9}$$

and for Design II is:

$$\tilde{Y}_2 = \max_{A_2 \in C,D,E} \hat{Y}_2 I(A_1 = A) + \max_{A_2 \in F,G,H} \hat{Y}_2 I(A_1 = B). \tag{3.10}$$

We then model to estimate the best first stage treatment:

$$\tilde{Y}_2 \sim \beta_A I(A_1 = A) + \beta_B I(A_1 = B)$$

The first stage treatment with the largest $\hat{\beta}$ is then selected.

3.4.2.3 Tree-based reinforcement learning

For the TBRL method with IPTW, we start similarly to Q-learning by using the appropriate weighted linear regression models from equations 3.5 and 3.6 or 3.7 and 3.8. We then calculate the 100-percentiles of the observed outcomes and test each as a potential cut point using the methods described in Section 3.3.3. The cut point that results in the highest $E(Y_2)$ based on the parameter estimates is selected if it meets all splitting criteria. For tree building, we follow Tao et al. (2018) and use $n_0 = 5$ as our minimum node size. We set $\lambda = 0.01$ as our minimum node purity improvement so that there is at least a 1% improvement in the predicted second stage outcome by allowing a split. Since we do not include covariates, we only have one possible variable for splitting (Y_1) and thus we do not need to implement Rule 3 regarding maximum tree-depth.

As with the Q-learning method, we estimate the best first stage treatment by maximizing psuedo-outcome, $\tilde{Y}_2$ as defined in 3.9 or 3.10. We select the first stage treatment with the largest $\hat{\beta}$.

3.4.2.4 Weighted and Replicated Regression

For the TV-SMART designs, we do not need to estimate the cut point as it is predefined. Additionally, it is uncommon for SMARTs using a tailoring variable to re-estimate the optimal cut off as it is assumed to be known. Therefore, we do not apply Q-learning or TBRL to the data from the TV-SMART design as this would not be done in practice.

For Design I, we only need to estimate two predefined DTRs that either start with treatment A or start with treatment B. Since patients are randomized to either A or B we only need to compare the second stage outcomes based on first stage treatment so the model is:

$$Y_2 \sim \beta_1 I(A_1 = A) + \beta_2 I(A_1 = B)$$

For Design II, we must select a second stage treatment for people whose Y_1 values fall below the pre-defined threshold. To do this, we use the weighting and replication method described in

Nahum-Shani et al. (2012a) and Kidwell et al. (2018). Here, we examine four possible DTRs: (A, C, D), (A, C, E), (B, F, G), and (B, F, H). Thus, for participants that receive treatment C, they could be following either of the first two DTRs, and for participants that receive treatment F, they could be following either of the second two DTRs. Thus we replicate those participants, one replicate for each DTR they could be following, to get an accurate estimate of the second stage outcome on each DTR. To account for imbalance by design, we weight participants using the weights in Equation 3.4. Let the selected cut off for the tailoring variable be x. Then $f(Y_1)$ in a TV-SMART is defined as:

$$f(Y_1) = \begin{cases} 0 & Y_1 \leq x \\ 1 & Y_1 > x. \end{cases}$$

We then use linear regression with IPTW to model:

$$Y_2 \sim \beta_1 I(DTR = (A, C, D)) + \beta_2 I(DTR = (A, C, E)) + \\ \beta_3 I(DTR = (B, F, G)) + \beta_4 I(DTR = (B, F, H))$$

The DTR is then selected based on which has the highest estimated β, $\hat{\beta}$.

3.4.3 Design Settings

We simulated 1000 trials that had a total of 300 patients with 150 patients on each of the first stage treatments, A and B.

3.4.3.1 Parameters for data generation models

For all designs and settings we set $\alpha_A = 50$, $\alpha_B = 30$, and all standard deviations to be $\sigma = 15$. We vary the second stage intercept, β, and slope, γ, parameters to create different scenarios (Tables 3.1 and 3.2). Scenarios I-i and II-i have the same cut off for the DTRs that start with either treatment A or treatment B. This follows most literature on estimating DTRs where there is not an interaction between A_1 and Y_1. Scenarios I-ii and II-ii have different cut offs for the DTRs depending on which first stage treatment is selected. In Scenarios I-iii and II-iii, $E(Y_2|A_2, A_1 = A) = E(Y_2|A_2, A_1 = B)$ meaning there is no difference between second stage outcomes for treatments $A_1 = A$ and $A_1 = B$ conditional on Y_1. There is still a difference in first stage outcomes so treatment A is the optimal first stage treatment. In scenarios I-iv and II-iv, the cut point is not within a reasonable range of Y_1 (within $\pm 3\sigma$ of the true cut off) so a DTR is not needed. Thus the optimal treatment sequence is (A, C) for all patients. For scenarios I-v and II-v, the optimal DTR is not clear from only examining the second stage expected outcome, $\mu(A_2, Y_1, A_1)$. In Scenario I-vi,

Y_1 is prognostic but not predictive, such that Y_1 is yields information about Y_2 but does not yield information about the association of A_2 with Y_2 conditional on A_1. In Scenario II-vi, the optimal DTR has two cut points, so there are three possible second stage treatment options as opposed to the traditional two second stage treatments.

3.4.3.2 Tailoring functions and tailoring variables

In our simulations, we consider three different tailoring functions. For all, if $Y_1 > 100$, $f(Y_1) = 1$ and if $Y_1 < 0$, $f(Y_1) = 0$. Then the three functions used are:

$$\text{TF 1: } f(Y_1) = Y_1/100$$
$$\text{TF 1/2: } f(Y_1) = (Y_1/100)^{1/2}$$
$$\text{TF 2: } f(Y_1) = (Y_1/100)^2$$

We also consider six different tailoring variables by assigning a cut point of the first stage outcome. If Y_1 is higher than the cut point, then $f(Y_1) = 1$ and if Y_1 is below the cut point then $f(Y_1) = 0$. The cut offs used in simulation are 30, 40, 50, 60, 70, 80. These are labelled as TV 30, TV 40, TV 50, TV 60, TV 70, and TV 80, respectively.

3.4.4 Performance metrics

3.4.4.1 Definition of the optimal DTR and optimal cut off

In order to compare results from the SMARTs to a defined standard, we use the optimal DTR defined for each scenario by the sets of treatments and cut offs listed in Tables 3.1 and 3.2. Since we present simulations with a known data generation model, we know the optimal DTRs, but this is not true in practice. For example, for Scenario I-i, the optimal DTR is to start all patients on treatment A. Then if the patient's Y_1 is greater than 60, they receive treatment C. If the patient's Y_1 is less than or equal to 60, they receive treatment D. In this scenario, the optimal cut off is 60. For scenarios that have only two treatments listed, (A, C), and "None" as the optimal DTR cut off for Y_1, the optimal DTR is to have everyone receive treatment A and then receive treatment C regardless of Y_1. There is no optimal cut off. For Scenario II-vi, the optimal DTR is to start everyone on treatment A. Then if $Y_1 > 60$, the patient receives C, if $40 < Y_1 \leq 60$ the patient receives D, and if $Y_1 \leq 40$ the patient receives E. This scenario has two optimal cut offs, 60 and 40. All optimal DTRs and associated optimal cut offs were calculated based on the known data generation parameters in Tables 3.1 and 3.2 as well as the known first stage treatment effects in Section 3.4.1.

Table 3.1: Parameters for each scenario explored for Design I in the data generation model for the second stage mean outcome shown in equation 3.2. DTR = dynamic treatment regimen, Y_1 = first stage outcome, β_{ij} = intercept term for when $A_1 = i$ and $A_2 = j$, γ_{ij} = slope term for when $A_1 = i$ and $A_2 = j$

Scenario	β_{AC}	γ_{AC}	β_{AD}	γ_{AD}	β_{BE}	γ_{BE}	β_{BF}	γ_{BF}	Optimal DTR	Optimal DTR cut off for Y_1
I-i	20	1.2	50	0.7	0	1.2	30	0.7	(A, C, D)	60
I-ii	20	1.2	50	0.7	10	1	30	0.3	(A, C, D)	60
I-iii	20	1.2	50	0.7	20	1.2	50	0.7	(A, C, D)	60
I-iv	40	1.2	20	1	20	1.2	0	1	(A, C)	None
I-v	20	1.2	50	0.7	10	2	0	1	(A, C, D)	60
I-vi	50	0.7	20	0.7	30	1	10	1	(A, C)	None

Table 3.2: Parameters for each scenario explored for Design II in the data generation model for the second stage mean outcome shown in equation 3.2. DTR = dynamic treatment regimen, Y_1 = first stage outcome, β_{ij} = intercept term for when $A_1 = i$ and $A_2 = j$, γ_{ij} = slope term for when $A_1 = i$ and $A_2 = j$

Scenario	β_{AC}	γ_{AC}	β_{AD}	γ_{AD}	β_{AE}	γ_{AE}	β_{BF}	γ_{BF}	β_{BG}	γ_{BG}	β_{BH}	γ_{BH}	Optimal DTR	Optimal DTR cut off for Y_1
II-i	20	1.2	50	0.7	30	0.3	0	1.2	30	0.7	10	0.3	(A, C, D)	60
II-ii	20	1.2	50	0.7	30	0.3	10	1.0	40	0.6	20	0.2	(A, C, D)	60
II-iii	20	1.2	50	0.7	30	0.3	20	1.2	50	0.7	30	0.3	(A, C, D)	60
II-iv	40	1.2	20	1	30	1.1	25	0.9	5	0.7	15	0.8	(A, C)	None
II-v	20	1.2	40	0.6	30	0.3	5	1.3	60	0.4	20	0.2	(A, C, D)	33.33
II-vi	10	1.2	40	0.7	60	0.7	20	0.6	40	0.3	30	0.1	(A, C, D, E)	60, 40

3.4.4.2 Definition of an estimated DTR

We estimate a DTR for each simulated trial data set using Q-learning and TBRL for TF-SMARTs and WRR for TV-SMARTs. For TF-SMARTs, the estimated DTR includes the treatments that make up the DTR and the estimation of the cut off for Y_1 (i.e. the definition of the tailoring variable). For TV-SMARTs, the estimated DTR includes the treatments that make up the DTR and the predefined cut off for Y_1 which is the same as the tailoring variable used in the trial.

3.4.4.3 Comparisons of SMART designs

For each design, scenario, and tailoring function/variable, we simulated 1000 trials and evaluate outcomes under each estimated DTR in sample of 10000 patients from the target population. We then estimated the DTR from one trial using the appropriate method(s) (Q-learning and TBRL for TF-SMARTs and WRR for TV-SMARTs). We apply resulting estimated DTR and apply it in a new population of 10,000 patients. We then average the second stage outcomes in the new population. We compare the average second stage outcomes, $\overline{Y}_2$, between TF-SMARTs and TV-SMARTs with higher being better. We also compare $\overline{Y}_2$ within TF-SMARTs for TF 1, TF 2, and TF 1/2 and within TV-SMARTs for TV 30, TV 40, TV 50, TV 60, TV 70, and TV 80 defined in Section 3.4.3.2. Additionally, we compare $\overline{Y}_2$ resulting from applying the estimated DTRs in the new population to the $\overline{Y}_2$ resulting from applying the optimal DTRs in the new population. We also examine the proportion of patients in the new population of 10,000 that are assigned the same treatments under the estimated DTRs as they would be under the optimal DTR. Lastly, we examine the proportion of trials that correctly selected treatment A as the optimal first stage treatment. In all scenarios, A is the optimal first stage treatment (Tables 3.1 and 3.2).

3.4.4.4 Comparison of analytic methods for TF-SMARTs

We compare Q-learning and TBRL as analytic methods for TF-SMARTs. To do this we examine the same metrics as described in Section 3.4.4.3: $\overline{Y}_2$, proportion of patients on the same treatment under the estimated DTR as they would be under the optimal DTR, and the proportion of trials that correctly selected treatment A as the optimal first stage treatment. For scenarios that had a single optimal cut off for Y_1 (Scenarios I-i, I-ii, I-iii, I-v, II-i, II-ii, II-iii, and II-v; Tables 3.1 and 3.2), we also examined how well the methods estimated the optimal DTR cut off for Y_1. To do this we looked at bias defined as the difference between the estimated cut off and the optimal DTR cut off. For Scenarios I-i, I-ii, I-iii, I-v, II-i, II-ii, and II-iii, the optimal DTR cut off is 60. For Scenario II-v, the optimal DTR cut off is 33.33.

3.5 Results

3.5.1 Comparison between tailoring functions and tailoring variables

The average second stage outcome, $\overline{Y}_2$, when applying the estimated DTRs from the TF-SMARTs to the a population of 10,000 was fairly consistent regardless of what tailoring function was used (Figure 3.6). However, there was significant variation in $\overline{Y}_2$ when using a pre-selected cut off tailoring variable in the TV-SMARTs. For most scenarios there was at least one cut off that performed similarly to the tailoring functions. As expected, the TV-SMARTs using the cut off of 60 as a tailoring variable (TV 60) performs well in most settings since that is the optimal cut off in scenarios I-i, I-ii, I-iii, I-v, II-i, II-ii, and II-iii. In these scenarios with the TV-SMART design using TV 60, 100% of patients in the new population end up on the same treatment as they would with the optimal DTR (Figure 3.7). Thus, these trials yield nearly equivalent performance to the optimal DTR. However, there were also cut offs that did not perform as well. Overall, there was a wider variation in $\overline{Y}_2$ when conducting a TV-SMART compared to a TF-SMART.

Interestingly, in Scenario II-v where there was not a clear DTR based on second stage models, there is extreme variance of the average second stage outcome based on which tailoring variable was used. This was due to the WRR method described in Section 3.4.2.4 selecting a DTR with the incorrect first stage treatment. In this scenario, the TV-SMART using TV 80 selected treatment B for the first stage treatment over treatment A 29.2% of the time. In Scenario I-v, the TF-SMART using TF 2 selected treatment B in 0.7% of simulations when using Q-learning and 0.8% when using TBRL. These results explain the outliers for the TF-SMART using TF 2 with Q-learning in Figure 3.6. In Scenario II-v, all tailoring function methods selected treatment B at least 0.3% of the time. All other design, settings, and methods selected treatment A 100% of the time. In trials that selected the incorrect first stage treatment, none of the patients then received the treatments recommended by the optimal DTR (Figure 3.7).

In scenarios where no DTR was useful (I-iv and II-iv), the TV-SMARTs had worse second stage outcomes than the TF-SMARTs. TF-SMARTs have the flexibility to select no DTR, whereas TV-SMARTs do not. Assigning any cut point for a tailoring variable had poor performance. Performance decreased as the cut off increased since more people were assigned to a different treatment. In Scenario II-iv where the optimal DTR was composed of three second stage treatment options as opposed to two, the TV-SMARTs never reached the same performance in $\overline{Y}_2$ as the TF-SMARTs which allow for the flexibility to identify three second stage treatment options. However, the percent of patients that receive the same treatments as would be assigned based on the optimal DTR is nearly equivalent between trials using TV 60 and the trials using tailoring functions.

For both Designs I and II, Scenarios i, ii, and iii had very similar results for the expected second stage outcome. This is because the optimal DTR including the optimal cut off point for Y_1 are

Figure 3.6: The average second stage outcome when applying the estimated DTR to a new sample of 10,000 patients from the target population. TBRL is done using a minimum purity improvement parameter of $\lambda = 0.01$ DTR = dynamic treatment regimen, Y_2 = second stage outcome, TBRL = tree-based reinforcement learning, WRR = weighted and replicated regression, TF = tailoring function, TV = tailoring variable

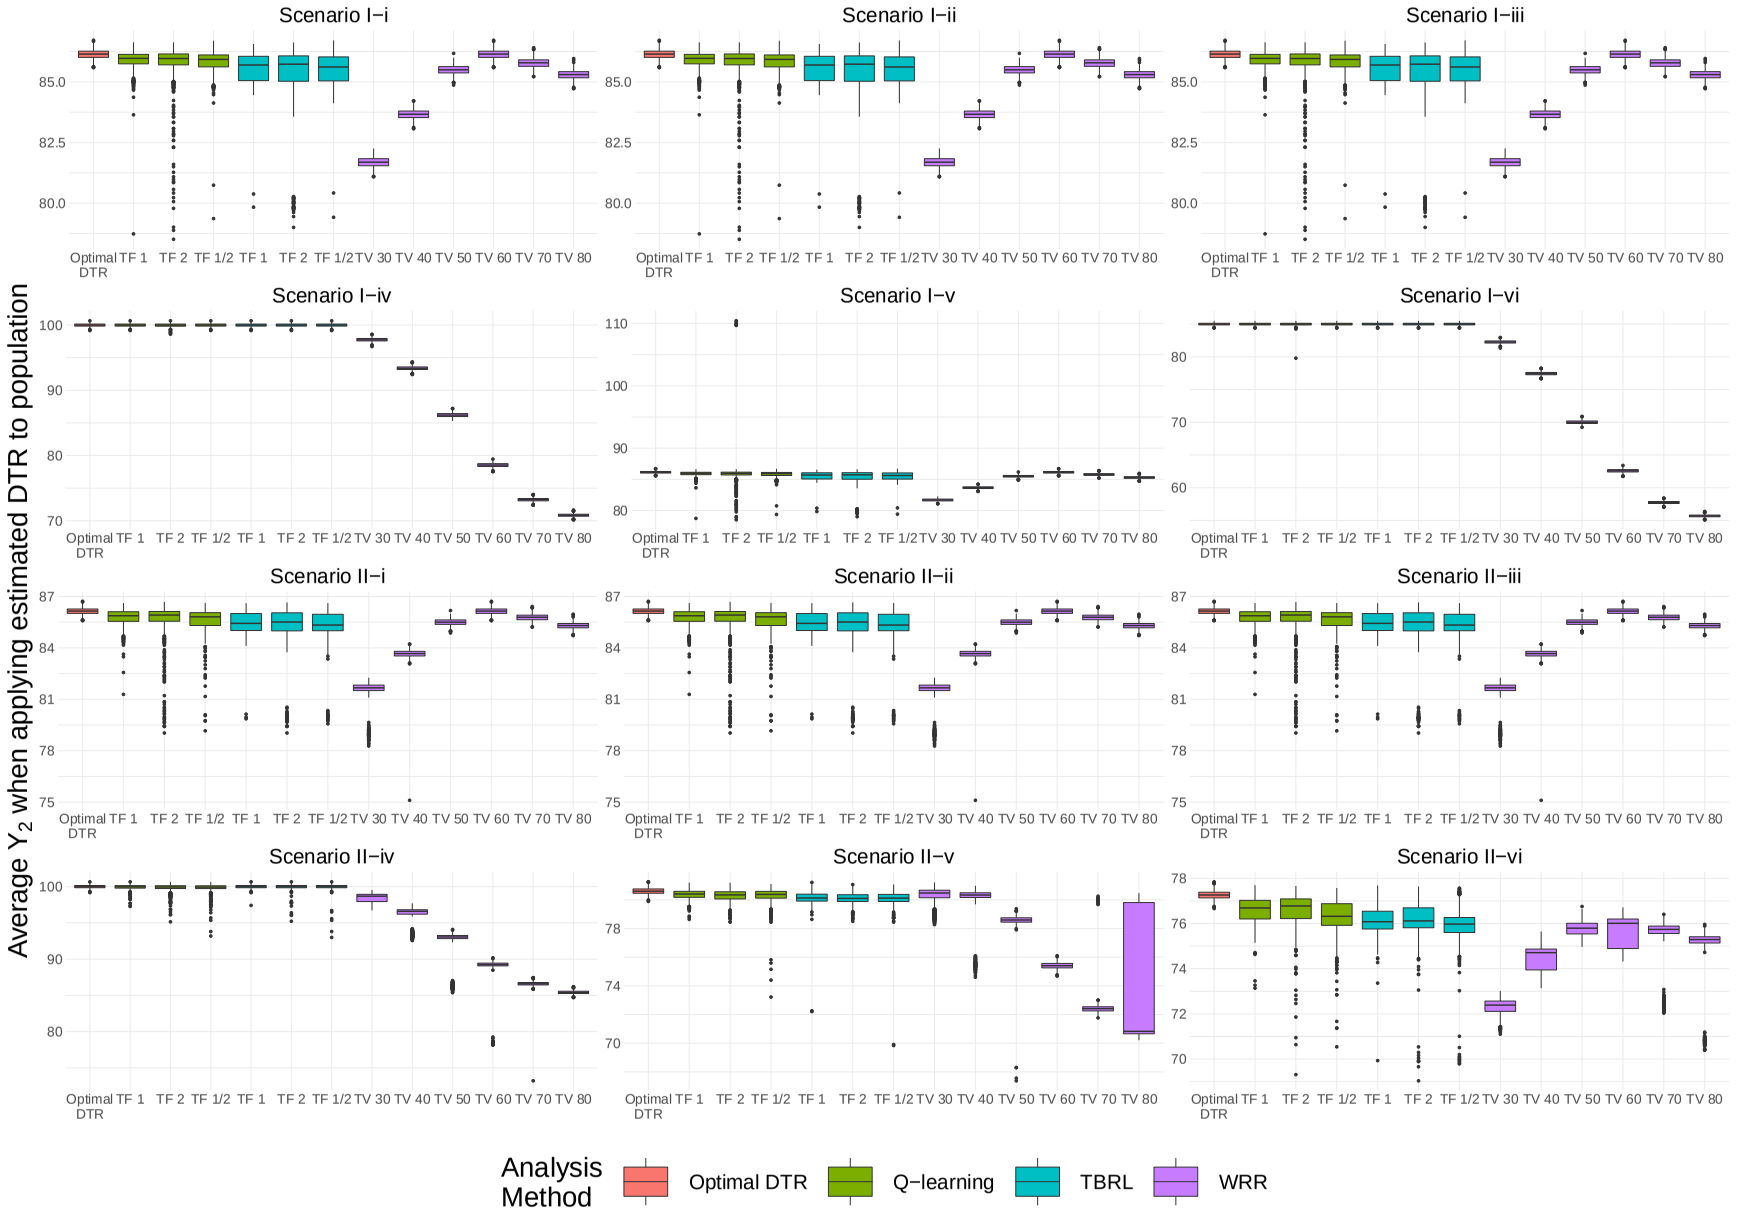

Figure 3.7: Overall proportion of patients in sample of 10,000 from the target population that end up on the same treatments as they would be assigned with the optimal DTR as defined in Section 3.4.4.1 and Tables 3.1 and 3.2. DTR = dynamic treatment regimen, TBRL = tree-based reinforcement learning, WRR = weighted and replicated regression, TF = tailoring function, TV = tailoring variable

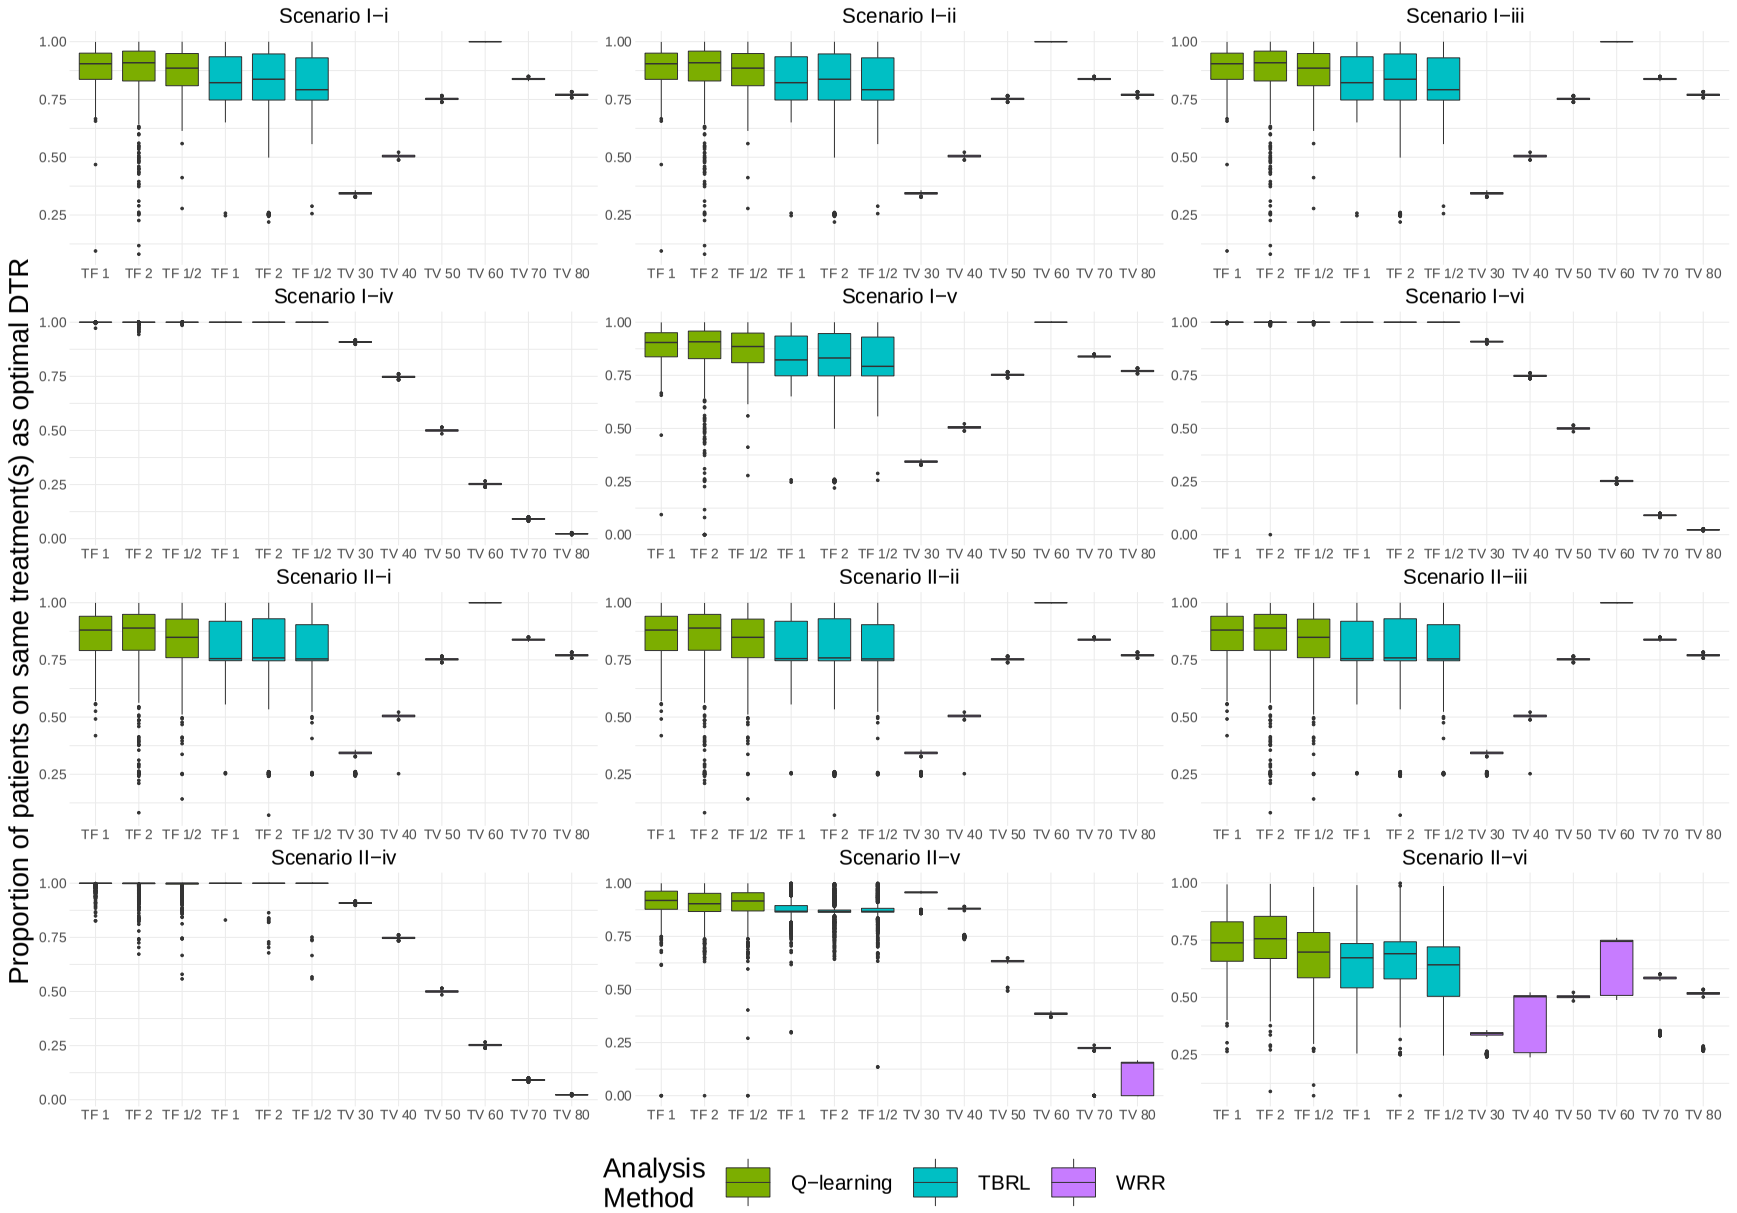

identical between these settings.

3.5.2 Comparison of analysis methods for TF-SMART designs

Q-learning tended to perform better than TBRL in terms of maximizing $\overline{Y}_2$ (Figure 3.6). The average second stage outcomes are higher with Q-learning, but there is also wider range of average second stage outcomes when using Q-learning. The performance of Q-learning varies greatly, especially when using TF 2. This is likely due to having very few patients that are randomized to certain treatment arms. The low number of patients in turn results in poor slope and intercept parameter estimates ($\hat{\beta}$s and $\hat{\gamma}$s, respectively) when modeling equations 3.5-3.8. Since the slope and intercept parameter estimates are directly used to estimate the DTR and cut off in Q-learning, more accurate estimation of the slope and intercept parameters yields better estimation of the cut off. Meanwhile, TBRL relies on the estimated parameters, splitting rules, and observed outcomes to define a cut off so it may be more robust to having fewer patients on certain treatment arms.

Next we wanted to compare the ability to estimate the correct optimal DTR cut off in a TF-SMART between Q-learning and TBRL by examining the bias as defined in Section 3.4.4.4. We only included simulations in which the estimated cut off was within $\pm 3\sigma$ of the optimal cut off as this is where the majority of first stage outcomes would lie based on the known distribution of Y_1. We also excluded cases when the TBRL did not make any splits based on the splitting rules defined in Section 3.3.3. The TBRL method did not select a cut off in over 40% of trials where the optimal DTR does include a cut off (Table 3.3). We found that analyzing the data using Q-learning resulted in less bias in the estimation of the cut off, but wider variance than using TBRL (Figure 3.8). For Q-learning, the median bias of the estimated cut off ranged between -0.754 to 0.856 across all scenarios. Meanwhile, for TBRL, the median bias of the estimated cut off ranged between -6.31 and 4.63 across all scenarios. The pattern where TBRL has larger bias than Q-learning also held in simulations where the trial size was 1000 patients with 500 on each initial arm (Figure 3.9) indicating that this bias is not due to insufficient patients on the trial to estimate the optimal split and cut off. When we changed the minimum purity requirement to allow a split from $\lambda = 0.01$ to $\lambda = 0.00001$, we saw improvement in the bias for this method (Figure 3.10). This also resulted in a smaller proportion of trials that incorrectly recommended no DTR (Table 3.4). This suggests that if the minimum improvement is set too high, the TBRL algorithm will have bias when estimating the optimal cut off. Additionally, when the minimum improvement parameter is set lower, the performance in terms of average second stage outcomes of the TBRL method becomes comparable to the performance of Q-learning (Figure 3.11). This suggests that this method requires tuning the minimum improvement parameter, λ, to achieve better performance. Despite the bias in the estimated DTR cut off when using TBRL with sample size of 300 and $\lambda = 0.01$, the majority

of patients still ended up on the same treatment as they would receive when using the optimal DTR (Figure 3.7).

Table 3.3: Proportion of trials that result in no DTR being selected such that all patients receive the same treatments when the total sample size is 300 and the minimum purity improvement $\lambda = 0.01$. Proportions should be 0 as all scenarios presented in this table have an optimal DTR and optimal cut off as defined in Section 3.4.4.1. DTR = dynamic treatment regimen, TBRL = tree-based reinforcement learning, TF = tailoring function

Analysis method	TF	I-i	I-ii	I-iii	I-v	II-i	II-ii	II-iii	II-v
Q-Learning	TF 1	0.021	0.021	0.021	0.021	0.052	0.052	0.052	0.025
Q-Learning	TF 2	0.025	0.025	0.025	0.025	0.048	0.048	0.048	0.078
Q-Learning	TF 1/2	0.041	0.041	0.041	0.041	0.088	0.088	0.088	0.053
TBRL	TF 1	0.431	0.431	0.431	0.431	0.441	0.441	0.441	0.582
TBRL	TF 2	0.403	0.403	0.403	0.403	0.435	0.435	0.435	0.583
TBRL	TF 1/2	0.419	0.419	0.419	0.419	0.453	0.453	0.453	0.578

Table 3.4: Proportion of trials analyzed using TBRL that result in no DTR being selected such that all patients receive the same treatments when the total sample size is 300 and the minimum purity improvement parameter is set to $\lambda = 0.00001$. Proportions should be 0 as all scenarios presented in this table have an optimal DTR and optimal cut off as defined in Section 3.4.4.1. DTR = dynamic treatment regimen, TBRL = tree-based reinforcement learning, TF = tailoring function

TF	I-i	I-ii	I-iii	I-v	II-i	II-ii	II-iii	II-v
TF 1	0.046	0.046	0.046	0.046	0.073	0.073	0.073	0.119
TF 2	0.056	0.056	0.056	0.056	0.086	0.086	0.086	0.202
TF 1/2	0.094	0.094	0.094	0.094	0.120	0.120	0.120	0.103

3.6 Discussion

Here we present a novel TF-SMART design that estimates the optimal DTR, including the cut off within the DTR, using a continuous intermediate outcome. This flexible design can result in better patient outcomes when applying the estimated DTR when compared to a TV-SMART design while circumventing the need to define a binary or categorical tailoring variable a priori. When the tailoring variable is a dichotomization of a continuous variable, there is wide variation in average second stage outcomes based on what cut off is used. When the cut off selected is not the true optimal cut off, there is a decrease in performance. Tailoring functions circumvent the need to select a cut off before the trial has been conducted. This allows researchers greater flexibility while maintaining good statistical properties.

Figure 3.8: The bias in the estimated cut offs relative to the optimal cut off with total sample size of 300 and the minimum purity measure $\lambda = 0.01$. The optimal cut off is defined in Section 3.4.4.1 and Tables 3.1 and 3.2. The bias is defined in Section 3.4.4.4. DTR = dynamic treatment regimen, TBRL = tree-based reinforcement learning, TF = tailoring function

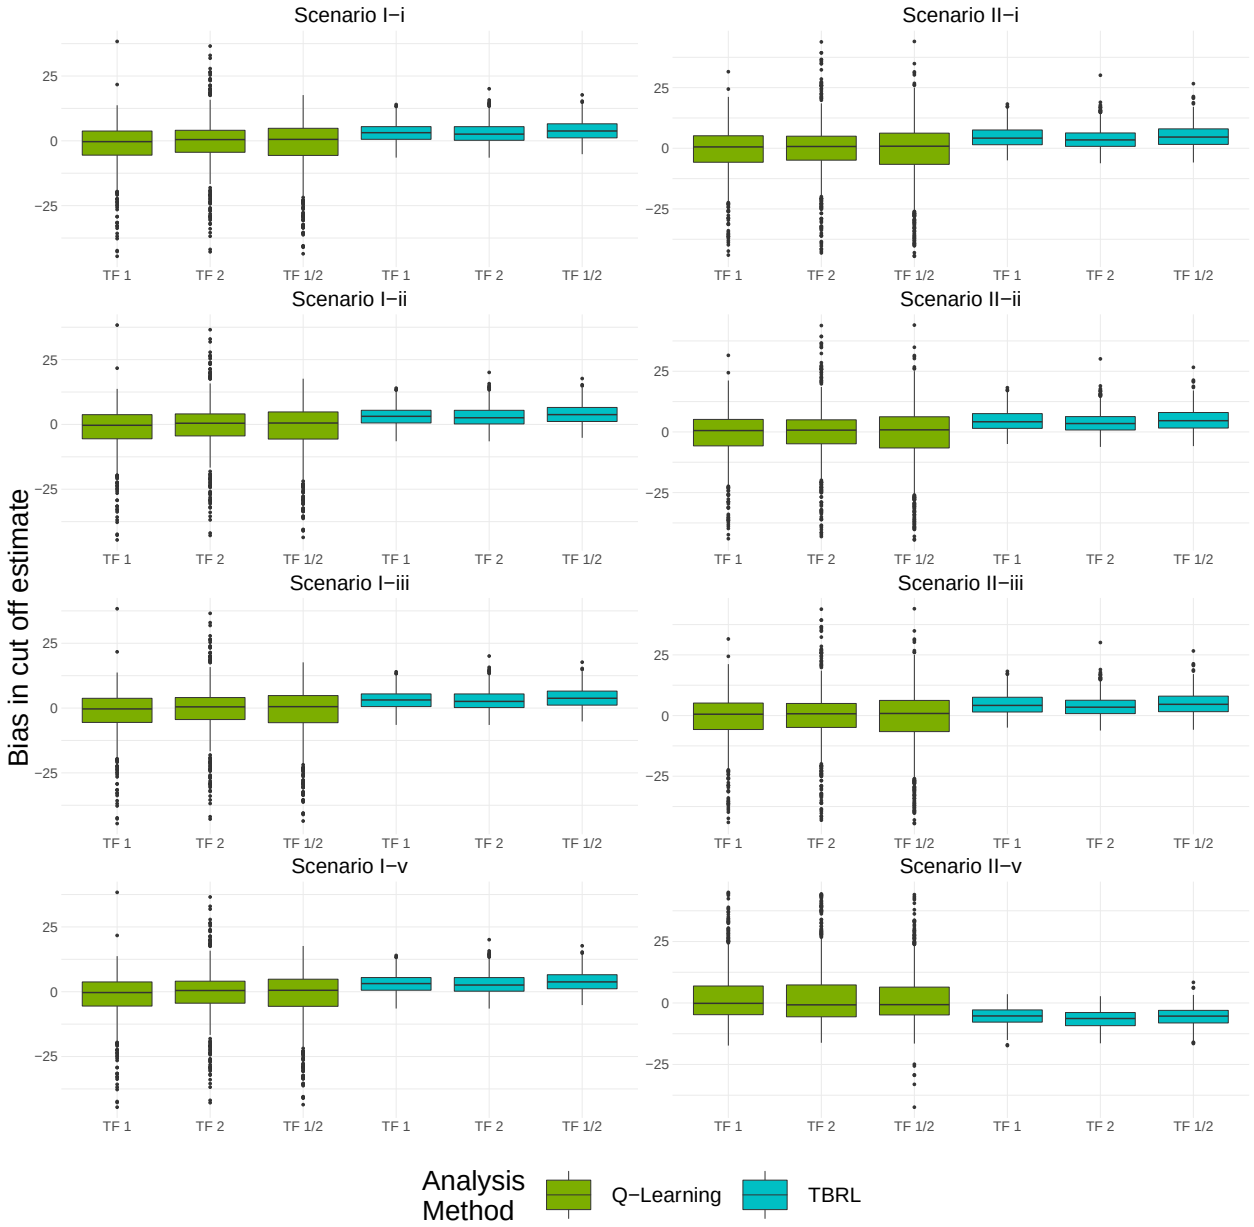

Figure 3.9: The bias in the estimated cut offs relative to the optimal cut off with total sample size of 1000 and the minimum purity measure $\lambda = 0.01$. The optimal cut off is defined in Section 3.4.4.1 and Tables 3.1 and 3.2. The bias is defined in Section 3.4.4.4. DTR = dynamic treatment regimen, TBRL = tree-based reinforcement learning, TF = tailoring function

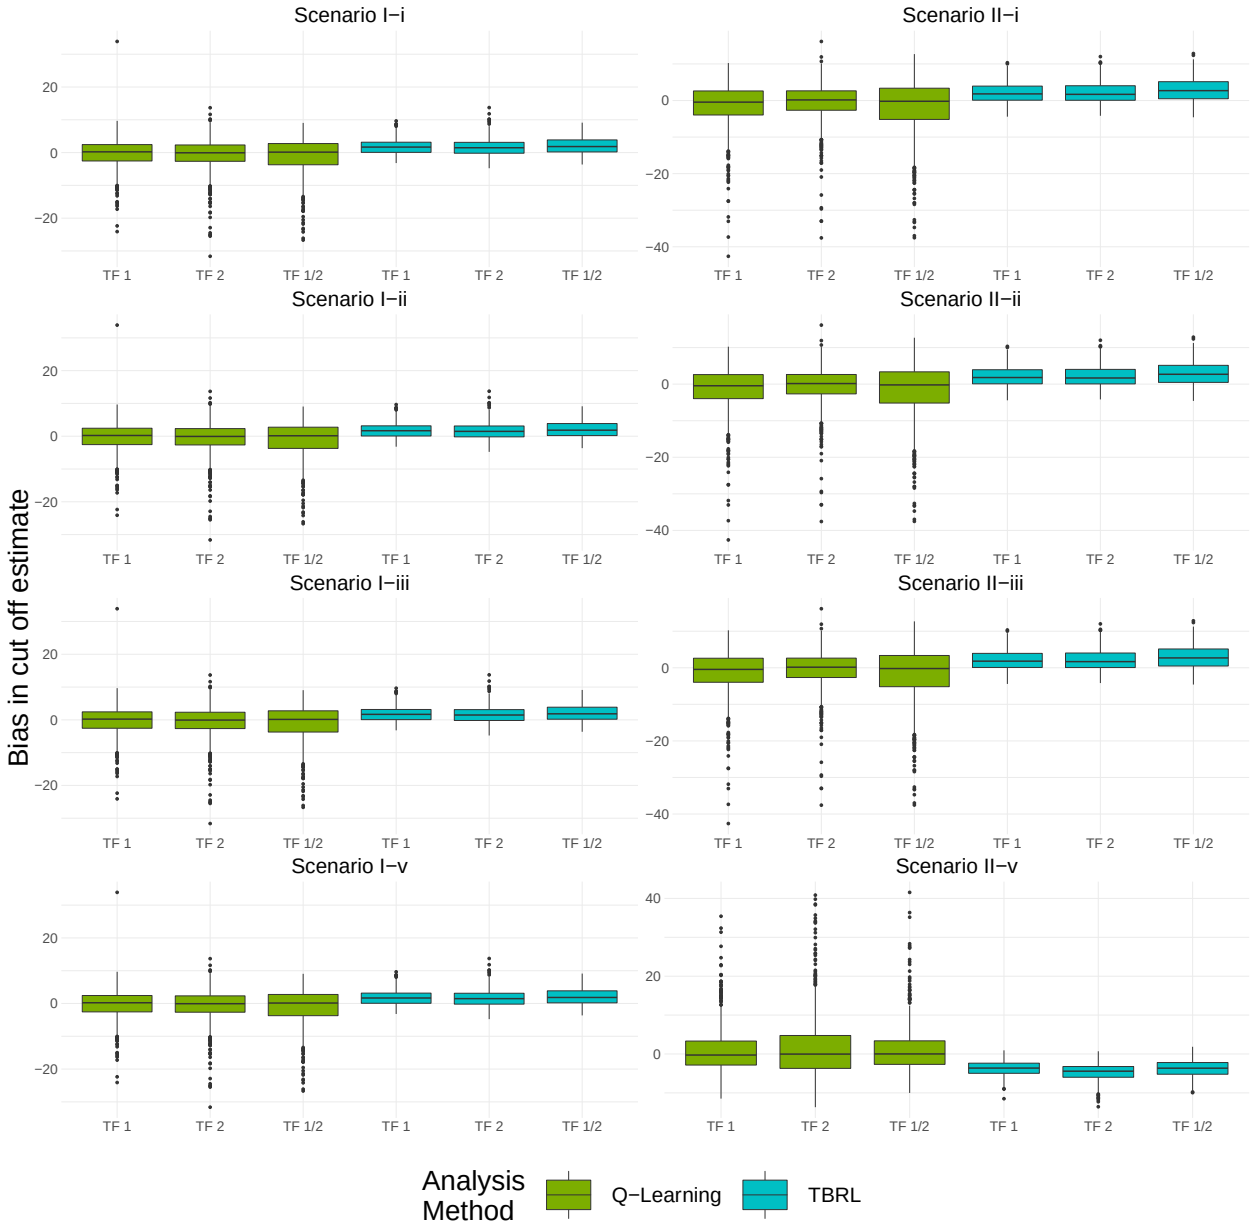

Figure 3.10: The bias in the estimated cut offs relative to the optimal cut off with total sample size of 300 and the minimum purity measure $\lambda = 0.00001$. The optimal cut off is defined in Section 3.4.4.1 and Tables 3.1 and 3.2. The bias is defined in Section 3.4.4.4. DTR = dynamic treatment regimen, TBRL = tree-based reinforcement learning, TF = tailoring function

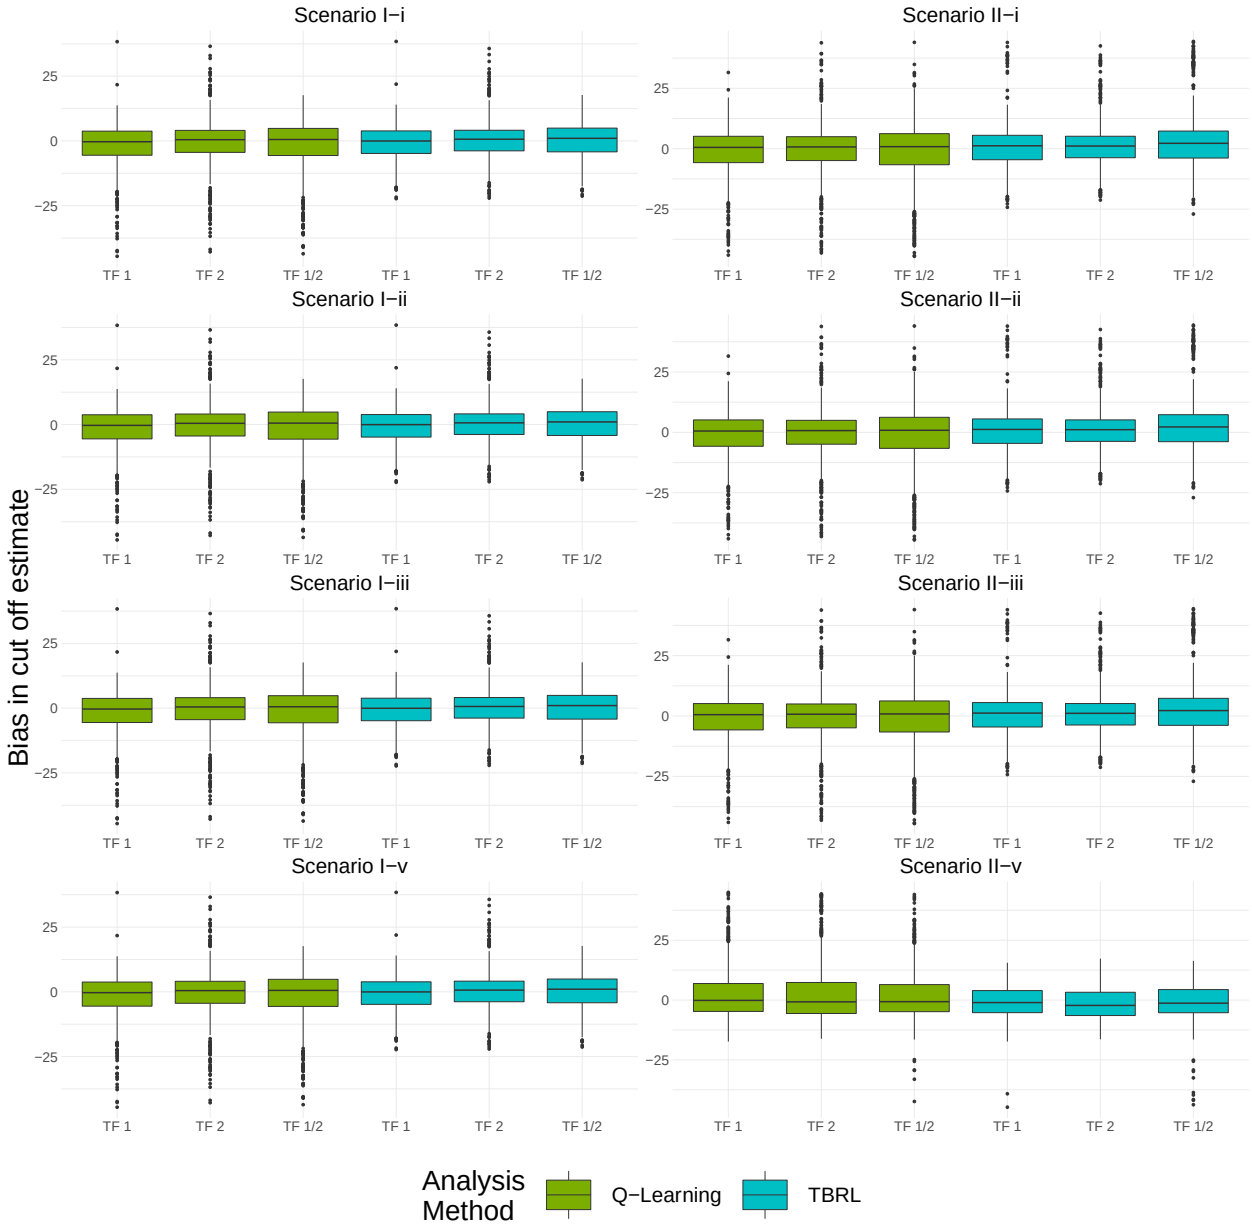

Figure 3.11: The average second stage outcome when applying the estimated DTR to a new sample of 10,000 patients from the target population. TBRL is done using a minimum purity improvement parameter of $\lambda = 0.00001$. DTR = dynamic treatment regimen, Y_2 = second stage outcome, TBRL = tree-based reinforcement learning, WRR = weighted and replicated regression, TF = tailoring function, TV = tailoring variable

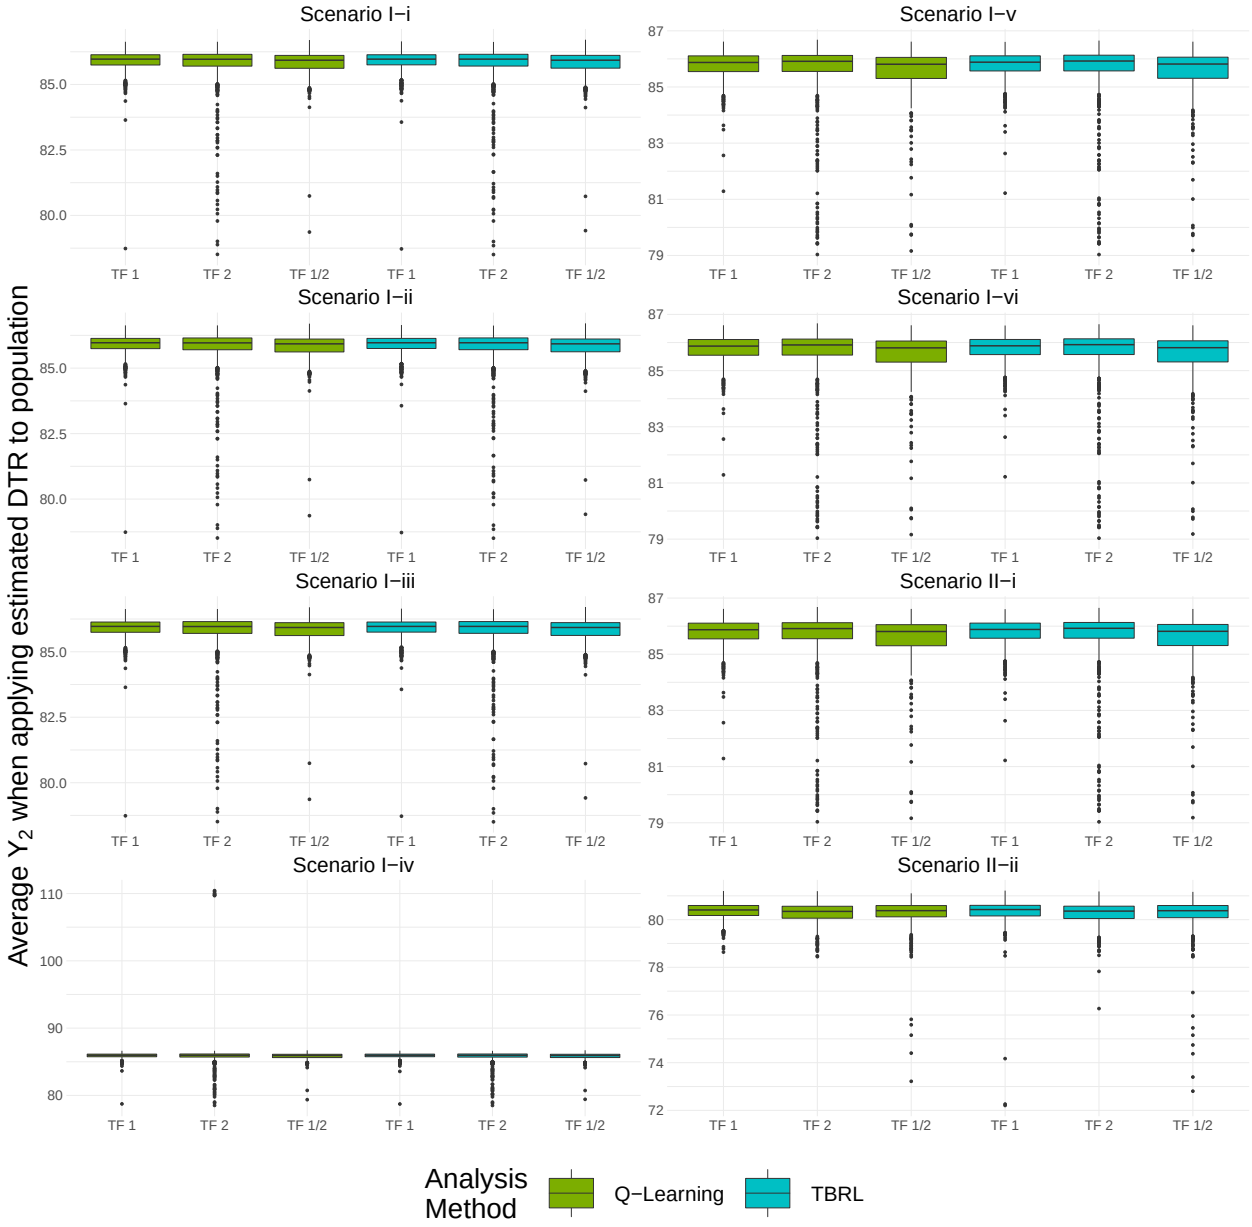

In our analysis methods, we focused on Q-learning and TBRL both use IPTW. However, there are many other methods that could be used as well as extensions to the methods presented here. There are methods that allow TBRL to estimate the minimum purity parameter improvement parameter in order to accurately predict when a split is warranted (Laber and Zhao, 2015; Tao et al., 2018). Since we saw a change in performance based on the minimum purity parameter, this may result in improved performance for this method. TBRL can also be implemented with random forests as opposed to linear regression (Tao et al., 2018). Additionally, there are other methods that have been used to estimate DTRs with observational data that could be implemented in the TF-SMART setting such as OWL (Zhao et al., 2012), sequential outcome weighted multicategory learning (Zhou et al., 2018), and Bayesian machine learning (Murray et al., 2018).

When randomization is appropriate, one major advantage to developing DTRs in a TF-SMART as opposed to using observational data is that propensity weights are known by design. We do not require a model to estimate the propensity weights. This is especially beneficial as estimates can be biased when the propensity model is misspecified (Drake, 1993; Lee et al., 2009; Rosenbaum and Rubin, 1983).

Although we presented models without covariates to illustrate the analysis of our trial design, it is easy to incorporate covariates. Both Q-learning and TBRL have been used to analyze observational data with a multitude of covariates to define DTRs and this can easily be translated to a randomized trial setting. The final product would be a DTR that could include covariates and the first stage outcome for a more personalized treatment regimen. While estimating sample size is outside the scope of this paper, incorporating covariates may require a larger sample size than a design that does not incorporate covariates. Simulations could be run to estimate the required sample size. In addition to including covariates in the tailoring variable for the estimated DTR, the tailoring function could also include covariates in the design stage. The tailoring function is flexible and could could be defined to accept more input than a single first stage outcome and still output a probability of receiving each second stage treatment.

In some trials, the difference in outcome between treatments and DTRs may be much smaller than what we presented here. If this is the case, then the sample size will need to be larger to be able to effectively identify the optimal DTR

Additional variations of the TF-SMART design beyond figures 3.1 and 3.2 are also feasible. The presented analytic methods can incorporate as many treatments and stages as practical. The methods here can also handle if there are uneven treatment arms where there are a different number of second stage treatment options based on first stage treatment. The weights in Section 3.4.2.1 and the models in Section 3.4.2.2 would need to be modified to match the design.

While the TF-SMART design presented here has the advantage that researchers do not need extensive knowledge about the optimal cut off a priori, if there is a wealth of knowledge, then

a traditional TV-SMART may be more appropriate. Alternatively, data from a TF-SMART design could be analyzed using Bayesian methods for defining a DTR (such as Bayesian Machine Learning as proposed by Murray et al. (2018)) with a strong prior to leverage that knowledge.

Overall, we demonstrated the use of a tailoring function in a SMART which offers a new and flexible way to develop DTRs in a clinical trial setting. By using a tailoring function to assign probabilities of receiving each treatment and applying methods for observational studies to our trial data, we expect to expand the scope of SMARTs and enable the development of effective DTRs across a wide variety of diseases and disorders.

CHAPTER 4

A Bayesian Approach to Comparing Overall Survival in Cancer Patients During the COVID-19 Pandemic Under Varying Treatment Timing Strategies

4.1 Introduction

When the COVID-19 pandemic hit, it interrupted all aspects of our society. One major dilemma facing oncologists was how to treat patients during the pandemic. Some hospitals had no choice but to close all departments to make room for the influx of COVID-19 patients. Other health care facilities (HCFs) were concerned that they would become overwhelmed with COVID-19 patients and needed to determine the optimal way to prioritize care for non-COVID-19 patients. Those HCFs that were able to remain open had to balance the risk of patients contracting COVID-19 by coming in for cancer treatment with the risk of worse cancer outcomes due to delaying treatment to minimize COVID-19 exposure. In addition, cancer patients may have higher mortality rates from COVID-19 (Landman et al., 2020; Williams et al., 2020) and vaccines may be less effective in immunocompromised patients such as those receiving chemotherapy (Center for Disease Control and Prevention, 2021). Simplistic tier systems were developed to determine if a cancer patient should receive cancer care as usual or if they should postpone their cancer treatment to reduce risk of contracting COVID-19 (American College of Surgeons, 2020; Association of Community Cancer Centers, 2020). These tier systems are limited since they do not take into consideration many patient and community factors that impact both cancer survival and COVID-19 infection risk and survival.

We previously developed a personalized model that predicts the overall survival of cancer patients during COVID-19 under both immediate cancer treatment and delayed cancer treatment. This project involved combining epidemiologic methods to estimate COVID-19 infection risk and

statistical methods for competing risks in survival analysis. In this tool we used a deterministic Susceptible-Infected-Recovered/Removed (SIR) model to estimate the current status of the pandemic in a patient's local region. We incorporated separate networks to allow for a higher rate of COVID-19 infection in HCFs. This assumption of higher risk in HCFs is supported by data showing that health care workers (HCWs) have a higher incidence rate than the general public (Center for Disease Control and Prevention, 2020; Nguyen et al., 2020). We used information about the individual patient's cancer (disease site, stage, patient age) to estimate overall survival in a pre-COVID-19 era. We also estimated the effect of cancer treatment delay on the patient's overall survival. We combined all of these estimates to develop estimates of overall survival for cancer patients during COVID-19. This model (OncCOVID) is published online at http://onccovid.med.umich.edu/ and in *JAMA Oncology* (Hartman et al., 2020). This model can be used worldwide for comparing treatment timing in patients with solid tumors.

In this manuscript, we further develop our previous methods and develop methodology to combine epidemic predictions of time-varying population infection and mortality rates as a competing risk to personalized survival estimates in a Bayesian framework. While much of the work involving pandemics and epidemics has been retrospective to examine properties of the disease and the population behavior, this work is prospective and integrates data that is changing daily to provide personalized patient level survival estimates that can be used in treatment decision making.

The methods presented here better capture the uncertainty and variance of our estimates. To do this, we use Bayesian methods to derive estimates for several quantities. First, we estimate the COVID-19 infection risk for both when the patient is receiving cancer treatment and when the patient is not. These estimates are based on the patient's county, the cancer treatment plan and timing, and the local epidemic properties of COVID-19. We use Bayesian extensions of the SIR model to develop these estimates. Next, we estimate the COVID-19 case fatality rate (CFR) for the patient given their age and number of comorbidities using a Bayesian logistic model. Mathematically, the CFR is the probability of death conditional on COVID-19 infection. To estimate the overall probability of dying from COVID-19 under a specific cancer treatment plan, we use the product of the estimates of the probability of COVID-19 infection and the CFR. We then estimate the patient's survival in the absence of COVID-19 using the Surveillance, Epidemiology, and End Results (SEER) data and perform a Bayesian survival analysis for each disease site, disease stage, and 5-year age group. To estimate the survival under a delay of cancer treatment, we estimate the hazard ratio (HR) for the effect of delaying cancer treatment using published literature and the National Cancer Database (NCDB) data. Finally, we combine the pre-COVID-19 overall survival estimate with the COVID-19 cause-specific survival estimate to derive a new overall survival estimate considering cancer treatment amidst a pandemic. Doing this within a Bayesian framework allows us to develop posterior predictive intervals for the estimated survival curves. Our primary

outcome of interest is the difference in restricted mean survival time (RMST) between immediate (treatment beginning 7 days after diagnosis) and delayed (treatment beginning 60 days after diagnosis) cancer treatment for an individual patient based on their specific covariates. These results can then be combined with physician's expertise and patient preferences to optimize the timing of cancer treatment thereby maximizing expected survival for the patient.

4.2 Methods

4.2.1 Overall procedure

Below is an outline of our procedure to estimate overall survival in cancer patients during the COVID-19 pandemic under delayed or immediate treatment. An overview of the methods, data sources, and resulting estimates are also shown in Figure 4.1.

1. Using the data from New York Times (2020) and models from the R package `EpiEstim`, simulate from posterior distribution for the current reproduction number (the average number of people infected from one infected person, R) for COVID-19 (Section 4.2.3.3).

2. Develop posterior distributions for the proportion of people currently infected with COVID-19 and the variance metric κ for the Dirichlet distribution using the Dirichlet beta state space model (Section 4.2.2.2; Osthus et al., 2017; Wang, Zhou, et al., 2020) and data from New York Times (2020).

3. Using a six state set up (susceptible, infected, and recovered for both the general public and HCWs) and a truncated Dirichlet method, estimate the posterior distribution of the COVID-19 epidemic curve based on the patient county's and our stated assumptions about COVID-19 (Section 4.2.3.1).

4. Estimate the COVID-19 CFR posterior distribution for each 10-year age group, sex, and number of comorbidities by conducting a Bayesian logistic regression using data from a small nested case-control study which is then standardized and normalized to data from a large cohort study (Section 4.2.4).

5. Estimate the posterior distribution of pre-COVID-19 survival for cancer patients using a piecewise constant hazard function model with a gamma process prior (Section 4.2.6.1).

6. Repeat the following 500 times for a given patient with a set treatment plan:

   (a) Estimate the posterior distribution of individualized COVID-19 specific survival curves:

i. Draw an epidemic curve from the posterior distribution from step 3.

ii. Based on the patient's cancer treatment plan (days in HCF) obtain the daily probability of becoming infected with COVID-19 using the drawn epidemic curve (Section 4.2.3.5).

iii. Draw a CFR from the posterior distribution in step 4 using the patient's age, sex, and number of comorbidities.

iv. Use the CFR draw and the daily probabilities of becoming infected to estimate the daily probability of dying from COVID-19 (Section 4.2.5).

(b) If considering delayed treatment, combine the pre-COVID-19 survival estimate with the appropriate HR (Section 4.2.6.2).

(c) Combine estimates of pre-COVID-19 survival and COVID-19 survival by multiplying survival estimates at time t for all t (Section 4.2.7).

4.2.2 Overview of methods to predict infectious disease spread

4.2.2.1 Deterministic Susceptible-Infected-Recovered Model

One commonly used method for estimating the trajectory of an epidemic or a pandemic is the SIR model (Cummings and Lessler, 2014). Let N be the population size, I_t be the number currently infected at time t, and R_t be the number recovered/removed. Then the number of people susceptible, S_t is $S_t = N - I_t - R_t$.

Let D is the duration of infectiousness and γ is $1/D$. β is the number of infectious contacts made during one day, also called the transmission rate.

The number of new infections on day $t+1$ is $\beta \frac{S_t I_t}{N}$. The number currently infected on day $t+1$ is $I_{t+1} = I_t + \beta \frac{S_t I_t}{N} - \gamma I_t$. The number susceptible on day $t+1$ is $S_{t+1} = S_t - \beta \frac{S_t I_t}{N}$. The number recovered/removed on day $t+1$ is $R_{t+1} = R_t + \gamma I_t$.

The SIR model (shown via a directed acyclic graph in Figure 4.2) can be viewed as the following set of ordinary differential equations:

$$\frac{dS}{dt} = -\beta SI \qquad \frac{dI}{dt} = -\beta SI - \gamma I \qquad \frac{dR}{dt} = \gamma I$$

R_0 is the number of other people infected from each infectious person, also called the reproductive number. This value is usually considered a constant property of an infectious disease based on unconstrained social interactions. However, through social distancing measures and as people become immune to COVID-19, the actual reproductive number can vary. The average number of people infected from one person on day t is R. Then $\beta = R/D$ (Cummings and Lessler, 2014).

Figure 4.1: A flow chart providing an overview of data sources (orange), statistical methodology (green), input patient data (light orange) and the output estimates (blue). CFR = case fatality rate, SEER = Surveillance, Epidemiology, and End Results, NCDB = National Cancer Database

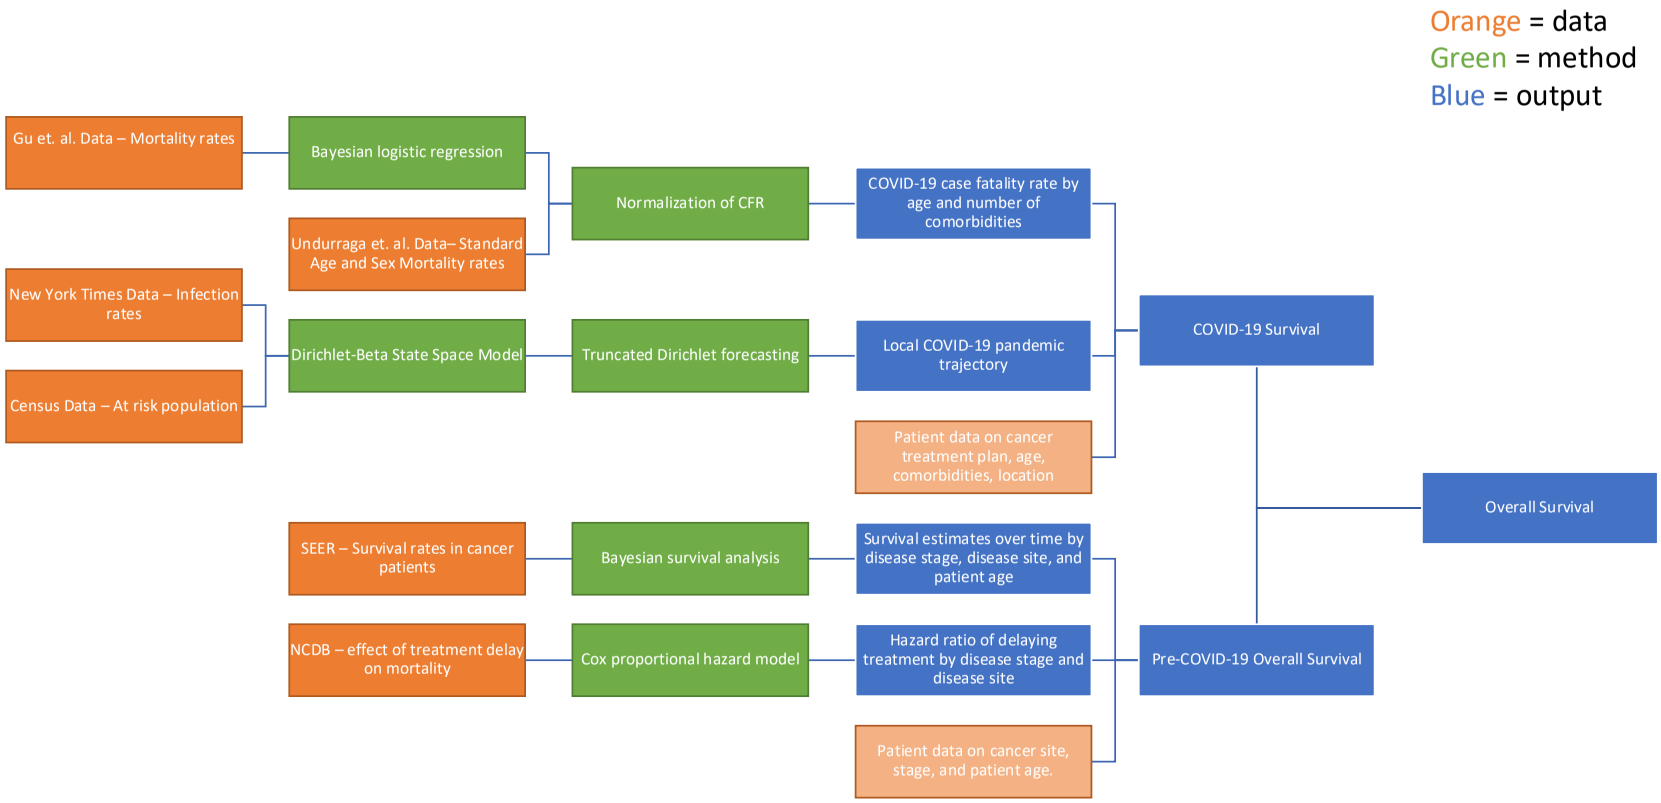

Figure 4.2: The susceptible-infected-recovered directed acyclic graph.

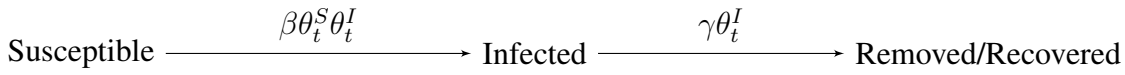

4.2.2.2 Dirichlet-Beta State-Space Model

This deterministic SIR was expanded to incorporate randomness in the Dirichlet-Beta state-space model (Osthus et al., 2017). In the Dirichlet-Beta state-space model, there are multiple states in which a person can be: susceptible, infected, or removed/recovered. The transition between these states is determined by a set of probabilities and distributions.

Let y_t be the proportion of the population that are a confirmed cases of COVID-19 at time t, $\theta_t = (\theta_t^S, \theta_t^I, \theta_t^R)'$ be the true proportion of susceptible, infectious, and recovered in the population, respectively. Let β and γ be as defined before. Let $\phi = \{\theta_0, \gamma, \beta, \kappa, \lambda\}$ where κ is a variance measure for the true proportion of people in each state and λ is a variance measure for the observed proportion of people in each state. In order to allow for uncertainty in the parameters, the disease transmission mechanism, and the case positive data, this method assumes that the proportion of confirmed cases follows a Beta distribution and the true underlying distribution of proportion of people in each state follows a Dirichlet distribution:

$$y_t|\theta_t, \phi \sim Beta(\lambda\theta_t^I, \lambda(1-\theta_t^I))$$
$$\theta_t|\theta_{t-1}, \phi \sim Dirichlet(\kappa f(\theta_{t-1}, \beta, \gamma))$$

From this, we see that $E(y_t) = \theta_t^I$ so the average proportion observed to be in the infected state on day t is equal to the true proportion of people in the infected state on day t. We define the function $f(\cdot)$ as the solution to the following ordinary differential equations:

$$\frac{d\theta_t^S}{dt} = -\beta\theta_t^S\theta_t^I \qquad \frac{d\theta_t^I}{dt} = -\beta\theta_t^S\theta_t^I - \gamma\theta_t^I \qquad \frac{d\theta_t^R}{dt} = \gamma\theta_t^I$$

Here $\frac{d\theta_t^S}{dt}$ is the change in the proportion of the population that is susceptible. This proportion is monotonically decreasing as people become infected. $\frac{d\theta_t^I}{dt}$ is the change in proportion of people that are currently infected. This increases when people become newly infected and decreases as people recover. $\frac{d\theta_t^R}{dt}$ is the change in the proportion of people that are recovered/deceased. This increases as people who were previously infected either recover from or succumb to the disease. Since a numerical solution is not known explicitly, $f(\cdot)$ is set to be equal to the fourth order Runge-Kutta approximation of the above equations (Osthus et al., 2017; Wang, Zhou, et al., 2020).

This method is used to produce posterior distributions for the number of people currently in-

fected with COVID-19, $\theta_1^{I_g}$, and the variance metric of the Dirichlet distribution, κ.

4.2.3 Estimating COVID-19 infection probability

4.2.3.1 Probabilistic state-space model with differential infection rate in mixing populations

As in a SIR model (Section 4.2.2.1; Cummings and Lessler, 2014), let R be the average number of other people infected from one infectious person, D be the duration of infectiousness, γ be $1/D$, and β be the number of infectious contacts made during one day (also called the transmission rate). Then $\beta = R/D = R\gamma$.

We use a Dirichlet-Beta state space model to derive the posterior distribution of the current proportion of the population that are susceptible to COVID-19, infected with COVID-19, or recovered/deceased from to COVID-19 (Section 4.2.2.2; Osthus et al., 2017; Wang, Zhou, et al., 2020).

For predicting the pandemic trajectory (i.e. forecasting), we use models that allow for two populations (HCWs and non-HCWs/general public) with differing transmission rates. The general public includes cancer patients in our context. The general public may go to a HCF and interact with HCWs. On days when the cancer patient is receiving treatment for their cancer, they are included as part of the general public that visits the HCF. Let h be proportion of the general public go to the HCF each day and H be the proportion of the overall population that are HCWs.

Let β_g be the transmission rate of COVID-19 in the general public and let β_H be the transmission rate at the HCF. Individuals that receive health care services on a given day have the risk of contracting COVID-19 both in the HCF and also in the general public. By adding two positive risks (risk in the HCF and risk in the general public), this forces a non-negative change in the probability of infection for the general public who go to the HCF which is appropriate given the current evidence that going to the HCF increases the risk of infection (Center for Disease Control and Prevention, 2020; Nguyen et al., 2020).

Let $\theta_t^g = (\theta_t^{S_g}, \theta_t^{I_g}, \theta_t^{R_g})$ be the proportion of the general public in each state on day t (susceptible: $\theta_t^{S_g}$, infected: $\theta_t^{I_g}$, recovered: $\theta_t^{R_g}$). Similarly, let $\theta_t^h = (\theta_t^{S_h}, \theta_t^{I_h}, \theta_t^{R_h})$ be the proportion of HCWs in each of the states.

Let $\theta_t^{S_H}$ be the proportion of susceptible people in the HCF. This includes the general public that at the HCF for care and HCWs that staff the HCF. Then this proportion, $\theta_t^{S_H}$, is the weighted average of the proportion susceptible in the general public and the proportion susceptible among HCWs as follows:

$$g_S = \frac{h}{h+H}\theta_t^{S_g},\ h_S = \frac{H}{h+H}\theta_t^{S_h},$$

$$\theta_t^{S_H} = g_S + h_S.$$

Similarly, for the proportion infectious at the HCF is expressed as follows:

$$\theta_t^{I_H} = \frac{h}{h+H}\theta_t^{I_g} + \frac{H}{h+H}\theta_t^{I_h}.$$

The mixing population Susceptible-Infected-Recovered/Removed (mpSIR) model (Figure 4.3) is equivalent to the following set of ordinary differential equations:

$$\begin{aligned}
\frac{d\theta_t^{S_g}}{dt} =& -h\left(\frac{g_S}{\theta_t^{S_H}}\beta_H\theta_t^{S_H}\theta_t^{I_H}\right) - (1-h)\left(\beta_g\theta_t^{S_g}\theta_t^{I_g}\right)\\
\frac{d\theta_t^{I_g}}{dt} =& h\left(\frac{g_S}{\theta_t^{S_H}}\beta_H\theta_t^{S_H}\theta_t^{I_H}\right) + (1-h)\left(\beta_g\theta_t^{S_g}\theta_t^{I_g}\right) - \gamma\theta_t^{I_g}\\
\frac{d\theta_t^{R_g}}{dt} =& \gamma\theta_t^{I_g}\\
\frac{d\theta_t^{S_h}}{dt} =& -\frac{h_S}{\theta_t^{S_H}}\beta_H\theta_t^{S_H}\theta_t^{I_H}\\
\frac{d\theta_t^{I_h}}{dt} =& \frac{h_S}{\theta_t^{S_H}}\beta_H\theta_t^{S_H}\theta_t^{I_H} - \gamma\theta_t^{I_h}\\
\frac{d\theta_t^{R_h}}{dt} =& \gamma\theta_t^{I_h}
\end{aligned} \tag{4.1}$$

Figure 4.3: The susceptible-infected-recovered directed acyclic graph where there is a separate infection risk when at health care facilities.

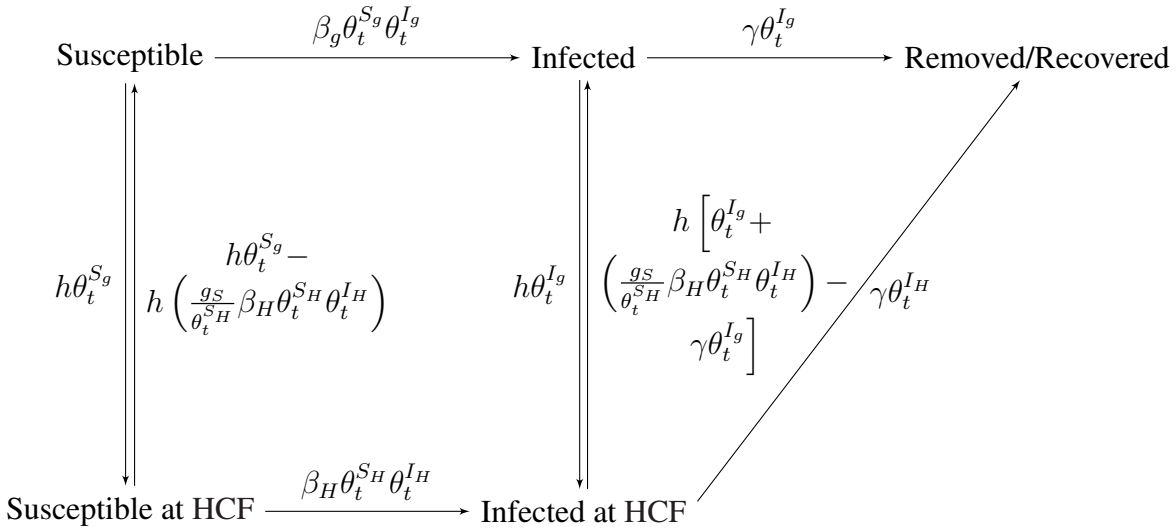

Here $\frac{d\theta_t^{S_g}}{dt}$ and $\frac{d\theta_t^{S_h}}{dt}$ are the change in the proportion of the general public and HCWs that are susceptible, respectively. $\frac{d\theta_t^{I_g}}{dt}$ and $\frac{d\theta_t^{I_h}}{dt}$ are the change in proportion of people that are currently

infected in the general public and HCW population, respectively. $\frac{d\theta_t^{R_g}}{dt}$ and $\frac{d\theta_t^{R_h}}{dt}$ are the change in proportion of people that are recovered/removed in the general public and HCW population, respectively.

The change in the number of people infected is entirely dependent on the change in the number of people who are susceptible and recovered. People who are no longer susceptible are infected and people who are no longer infected are recovered. Thus, the change in the number of people infected is equal to the number of people that are no longer susceptible minus the number of people that are newly recovered. Then:

$$\frac{d\theta_t^{I_g}}{dt} = \frac{d\theta_t^{S_g}}{dt} - \frac{d\theta_t^{R_g}}{dt} \text{ and}$$
$$\frac{d\theta_t^{I_h}}{dt} = \frac{d\theta_t^{S_h}}{dt} - \frac{d\theta_t^{R_h}}{dt}.$$

4.2.3.2 Truncated Dirichlet Distribution for Forecasting

One issue with the Dirichlet-Beta state space model for forecasting is that it allows for non-monotonicity in the true proportions of susceptible and recovered groups (i.e. the proportions in these groups can both increase and decrease) which is impossible under our assumptions. If violated, this would mean that people who have contracted COVID-19 can be reinfected which has been shown to be rare (Hall et al., 2021). Thus, we use truncated Dirichlet distributions for predicting the future trajectory of the pandemic. To draw from this truncated Dirichlet distribution we used methods described in Fang et al. (2000). Here we present an overview of this method, and more details can be found in Fang et al. (2000).

Let $\mathbf{x} = (x_1, \cdots, x_n)^T \sim TD(\mathbf{g}, \mathbf{h}; \gamma_1, \cdots, \gamma_n)$ be a truncated Dirichlet distribution. Here $\mathbf{g}$ and $\mathbf{h}$ are vectors of length n indicating the minimum values and maximum values respectively and $\gamma_1, \cdots, \gamma_n$ are the parameters of the Dirichlet distribution.

To draw from this truncated Dirichlet distribution we used methods outlined in Fang et al. (2000). Let $TBeta(g, h; \gamma_1, \gamma_2)$ be a truncated Beta distribution with shape and scale parameters γ_1 and γ_2, respectively, a minimum value of g, and maximum value h. Let $u_1, \cdots, u_{n-1}$ be distributed i.i.d $Uniform(0, 1)$. Fang proposes that if $\mathbf{x} = (x_1, \cdots, x_n)^T \sim TD(\mathbf{g}, \mathbf{h}; \gamma_1, \cdots, \gamma_n)$ then $\mathbf{x}$ has

the properties that:

$$
\begin{aligned}
x_{n-1} &\stackrel{d}{=} F_{n-1}^{-1}(u_{n-1}), \\
x_k &\stackrel{d}{=} \left(1 - \sum_{j=k+1}^{n-1} x_j\right) F_k^{-1}(u_k), k = n-2, n-3, \cdots 1, \\
x_n &\stackrel{d}{=} 1 - \sum_{k=1}^{n-1} x_k,
\end{aligned}
$$

where $\stackrel{d}{=}$ indicates the quantities have the same distribution, F_k^{-1} is the inverse of the truncated Beta distribution:

$$TBeta(\xi_k, \eta_k; \gamma_k, \sum_{i=1}^{n} \gamma_i - \sum_{j=k}^{n-1} \gamma_j)$$

where

$$\xi_k = \max\left(\frac{g_k}{1 - \sum_{j=k+1}^{n-1} x_j}, 1 - \frac{\sum_{i=1}^{n} h_i - \sum_{j=k}^{n-1} h_j}{1 - \sum_{j=k+1}^{n-1} x_j}\right)$$

and

$$\eta_k = \max\left(\frac{h_k}{1 - \sum_{j=k+1}^{n-1} x_j}, 1 - \frac{\sum_{i=1}^{n} g_i - \sum_{j=k}^{n-1} g_j}{1 - \sum_{j=k+1}^{n-1} x_j}\right).$$

We use the same relationships defined in Section 4.2.3.1 equation 4.1 and assign the estimated differential equations as the parameters of the truncated Dirichlet distributions.

$$
\begin{aligned}
\theta_t^g &\sim TD((0, 0, \theta_{t-1}^{R_g}), (\theta_{t-1}^{S_g}, 1, 1); \kappa(\theta_{t-1}^{S_g} + \frac{d\theta_t^{S_g}}{dt}, \theta_{t-1}^{I_g} + \frac{d\theta_t^{I_g}}{dt}, \theta_{t-1}^{R_g} + \frac{d\theta_t^{R_g}}{dt}) \\
\theta_t^h &\sim TD((0, 0, \theta_{t-1}^{R_h}), (\theta_{t-1}^{S_h}, 1, 1); \kappa(\theta_{t-1}^{S_h} + \frac{d\theta_t^{S_h}}{dt}, \theta_{t-1}^{I_h} + \frac{d\theta_t^{I_h}}{dt}, \theta_{t-1}^{R_h} + \frac{d\theta_t^{R_h}}{dt})
\end{aligned}
$$

Thus, using the results from Fang et al. (2000), the distributions for the proportion of the general public in each of the states of recovered, susceptible, and infected can be written as transformations of inverse truncated Beta distributions.

$$
\begin{aligned}
\theta_t^{R_g} &\stackrel{d}{=} F_{R_g}^{-1}(u_1), \\
\theta_t^{S_g} &\stackrel{d}{=} (1 - \theta_t^{R_g}) F_{S_g}^{-1}(u_2), \\
\theta_t^{I_g} &\stackrel{d}{=} 1 - \theta_t^{S_g} - \theta_t^{R_g}
\end{aligned}
$$

where $\stackrel{d}{=}$ indicates the quantities have the same distribution, u_1 and u_2 are independent

$Uniform(0,1)$ variables, $F_{R_g}^{-1}$ is the inverse of the truncated Beta distribution:

$$TBeta(\theta_{t-1}^{R_g}, 1; \kappa(\theta_{t-1}^{R_g} + \frac{d\theta_t^{R_g}}{dt}), \kappa(\theta_{t-1}^{I_g} + \frac{d\theta_t^{I_g}}{dt} + \theta_{t-1}^{S_g} + \frac{d\theta_t^{S_g}}{dt})),$$

and $F_{S_g}^{-1}$ is the inverse of the truncated Beta distribution:

$$TBeta(0, \frac{\theta_{t-1}^{S_g}}{1-\theta_t^{R_g}}; \kappa(\theta_{t-1}^{S_g} + \frac{d\theta_t^{S_g}}{dt}), \kappa(\theta_{t-1}^{I_g} + \frac{d\theta_t^{I_g}}{dt}))$$

4.2.3.3 Compilation of required estimates and published COVID-19 literature

In order to estimate the trajectory of the pandemic in a patient's county, we need estimates of the duration of infectiousness, D, and the time to recovery or death due to COVID-19, γ. The estimates of D and γ can be obtained from published data or literature.

To estimate the reproduction number R, we used COVID-19 case data over time from the New York Times (New York Times, 2020) and implemented the package `EpiEstim` in R (Cori et al., 2013). US county population size was obtained from 2019 US census estimates (US Census, 2020) and used to estimate the proportion of people in each state at $t = 0$. With these values, we can estimate the transmission rate in the general public, β_g.

4.2.3.4 Prediction method validation

We compared predictions using our mpSIR method with predictions from the Institute for Health Metrics and Evaluation (2020) for that same date, and the true number of confirmed cases according to the New York Times data (New York Times, 2020). Since Institute or Health Metrics and Evaluation (IHME) does not provide county level predictions, we applied our method to whole states. For our mpSIR method, we used data up until a pre-defined historic start date and then made predictions for 90 days after that date. We selected start dates where IHME had published predictions to ensure fair comparisons. The start dates selected were May 5, 2020; August 6, 2020; November 12, 2020; and February 4, 2021. We accessed the confirmed cases from the New York Times (denoted "truth" or "true confirmed cases", although we acknowledge that this data may be underreporting or have other errors) up to April 14, 2021.

4.2.3.5 Estimating an individuals risk of contracting COVID-19 based on treatment plan

The proportion of new infections in the general public that are due to exposure in the HCF on day t is:

$$N_t = \frac{h\left(\frac{h}{h+H}\beta_H \theta_t^{S_H}\theta_t^{I_H}\right)}{h\left(\frac{h}{h+H}\beta_H \theta_t^{S_H}\theta_t^{I_H}\right) + (1-h)\left(\beta_g \theta_t^{S_g}\theta_t^{I_g}\right)}.$$

Let P_t indicate if the patient is going into the HCF on day t and let $-\widehat{\frac{d\theta_t^{S_g}}{dt}}$ be the predicted proportion newly infected on day t using the truncated Dirichlet method in Section 4.2.3.2. Then the patient's individual risk of contracting COVID-19 on day t is:

$$F_t = -(1-N_t)\widehat{\frac{d\theta_t^{S_g}}{dt}} - P_t N_t \widehat{\frac{d\theta_t^{S_g}}{dt}}.$$

4.2.4 Estimating the case fatality rate for COVID-19 by age and number of comorbidities

COVID-19 mortality by age and number of comorbidities are estimated from data collected for a nested case control study (Gu et al., 2020) using Bayesian constrained logistic regression. The data are not restricted to cancer patients, but cancer is one of the comorbidities considered. The other comorbidities considered were hypertension, coronary heart disease, cardiac failure, cerebral infarction, diabetes, chronic bronchitis, chronic obstructive pulmonary disease, renal failure, and history of surgery.

The age groups are: $\leq$ 50, 50-60, 60-70, 70-80, and 80+. The number of comorbidities are: 0, 1, 2, and 3+. Let a be the person's age and c be the person's number of comorbidities. Letting Y be death due to COVID-19, then the model is:

$$\begin{aligned} logit\{P(Y=1|a,c)\} = \beta_0 + \beta_1 I(a \geq 50) + \beta_2 I(a \geq 60) + \beta_3 I(a \geq 70) + \\ \beta_4 I(a \geq 80) + \beta_5 I(c \geq 1) + \beta_6 I(c \geq 2) + \beta_7 I(c \geq 3) \end{aligned}$$

We constrain β_1 to β_7 to be non-negative by setting the priors to be folded normal distributions with a mean of 0 and a variance of 1000. The prior for β_0 is a normal distribution with mean 0 and variance of 1000. Since the data are from a case control study, we adjust the intercept for the above model to obtain accurate β parameter estimates. Let $\hat{\beta}_0$ be the estimated intercept from the model using data from Gu et al. (2020), and let $\hat{\pi}$ be the known proportion of patients that died in the Gu et al. (2020) study. Let $\tilde{\pi}$ be the estimated population CFR from The Novel Coronavirus Pneumonia Emergency Response Epidemiology Team (2020). Then the new, adjusted intercept is:

$$\tilde{\beta}_0 = \hat{\beta}_0 + \log\left(\frac{\tilde{\pi}}{1-\tilde{\pi}}\right) - \log\left(\frac{\hat{\pi}}{1-\hat{\pi}}\right).$$

The estimated log odds of dying from a COVID-19 infection for a person of age a and c comorbidities is:

$$\hat{B}_{ac} = \widehat{logit}(Y|a,c) = \tilde{\beta}_0 + \hat{\beta}_1 I(a \geq 50) + \hat{\beta}_2 I(a \geq 60) + \hat{\beta}_3 I(a \geq 70) + \\ \hat{\beta}_4 I(a \geq 80) + \hat{\beta}_5 I(c \geq 1) + \hat{\beta}_6 I(c \geq 2) + \hat{\beta}_7 I(c \geq 3),$$

and the case fatality rate estimate is:

$$\hat{E}_{ac} = \frac{\exp\left(\hat{B}_{ac}\right)}{1 + \exp\left(\hat{B}_{ac}\right)}.$$

Since the sample used to get these estimates consisted of 275 patients, we used a larger study (444,921 patients) to standardize the estimates so that the overall age CFRs were equal to the results from the larger study (Undurraga et al., 2021). This study provided case fatality rates by age group and sex. It did not provide age and comorbidity case fatality rates jointly. We also did not have access to this raw data and only summary statistics were published. Thus, to standardize, we obtained estimates of the percent of individuals in each age group with varying numbers of comorbidities from The National Health and Nutrition Examination Survey (NHANES) (Centers for Disease Control and Prevention, 2020). Data from the 2005-2006 cycle was used, and restricted to patients ages greater than 40 years old with complete information on age and sex and at least one non-missing entry for the comorbidities we considered. This yielded a final sample size of 3,056 US adults. Age specific prevalence values of having 0 to $\geq$3 comorbidities were used to calculate weighted average CFRs for each age decade. This results in age group CFRs that match the larger study, but maintain the relationships between the age and comorbidity levels seen in the Gu et al. (2020) data. We removed any estimates that resulted in probabilities greater than 1 after normalization.

4.2.5 Combining COVID-19 infection risk and COVID-19 case fatality rate to develop COVID-19 specific survival estimates

Let M_1 be the patient's COVID-19 CFR before cancer treatment and M_2 be the patient's CFR after cancer treatment begins. Note that $M_1 = M_2$ if the patient is not receiving chemotherapy. The immunosuppressive effect of chemotherapy is assumed to last 28 days after the end of treatment (Breastcancer.org, 2014). Let τ indicate the day when treatment starts and ω be when the treatment ends. We assume that when death occurs due to COVID-19 that it occurs 18 days after infection (Zhou et al., 2020). Thus, the mortality risk on day t depends on the probability of infection on day

$t-18$. Letting F_t be the individual risk of contracting COVID-19 on day t as defined in Section 4.2.3.5 and $t_{18} = t-18$, the COVID-19 mortality risk for day t is:

$$M_{COVID}(t) = [I(t_{18} < \tau \text{ or } t_{18} \geq \omega + 28)M_1 + I(t_{18} \geq \tau \text{ and } t_{18} < \omega + 28)M_2]F_{t_{18}}.$$

The probability $M_{COVID}(t)$ is absolute and is not conditional on survival up to time t. Thus cumulative probability of mortality over time T is additive rather than multiplicative and is calculated as:

$$\sum_{t=1}^{T} M_{COVID}(t)$$

We assume that $F_{t-18} = 0$ for $t < 18$. Then the COVID-19 specific survival is at time T for a patient with covariates x is:

$$S^{COVID}(t|x) = 1 - \sum_{t=1}^{T} M_{COVID}(t).$$

4.2.6 Estimating pre-COVID-19 overall mortality

Next, we estimated survival in cancer patients pre-COVID-19 using SEER data. Patients with invasive cancer were abstracted from the National Cancer Institute's SEER program. SEER collects data from 28% of the US population via a network of population-based incident tumor registries from geographically distinct regions in the US. The SEER 18 registry was used and all patients with an invasive cancer diagnosis were included from years 2005-2006. Patients with hematologic malignancies and patients diagnosed with cancer via autopsy or death certificate (<1.5% of patients) were excluded. The years of diagnoses were chosen in order to be both representative of contemporary cancer patients diagnosed in the US, and have the ability to generate long-term survival estimates post-diagnosis. A total of 691,854 patients met eligibility criteria with up to 12 years of follow up. Mortality codes in SEER are assigned from death certificates, completed by the doctor caring for the patient at the time of demise. From this dataset, 25 cancer types were extracted.

4.2.6.1 Estimating pre-COVID-19 mortality under standard of care

Using SEER data, we developed survival estimates for cancer patients and posterior predictive intervals by age, stage, and cancer type. We conducted a Bayesian survival analysis using a gamma process prior with a piecewise constant baseline hazard out to 10 years (Ibrahim et al., 2001; Sinha and Dey, 1997). We subset the SEER data by age group (5 year intervals), stage, and cancer type and estimated posterior survival curves for each subset separately.

This process is described further in Ibrahim et al. (2001), however the details of our process for developing a posterior survival curves are also described here. First, we estimate the posterior cumulative hazard. To do this, we set intervals equal to the observed event times, $s_1, s_2, ..., s_J$. Then we compute the number at risk (r_j) and the number of events/deaths (d_j) for each of the $j \in J$ time intervals. Let c_0 to be a weight (or mass) for the gamma process prior. We set the cumulative baseline hazard to be equal to the empirical Bayes estimate:

$$\lambda = \frac{\text{Total number of deaths}}{\text{Total observation time}}.$$

Then the estimated hazard at each time interval is:

$$\alpha_j = c_0 \lambda s_j.$$

The approximate posterior distribution for the increment in baseline hazard for the jth interval is (Ibrahim et al., 2001):

$$h_j \sim G(\alpha_j - \alpha_{j-1} + d_j, c_0 + r_j + d_j).$$

The estimated pre-COVID-19 survival at time t for a patient with covariates x is calculated as:

$$S^{pre}(t|x) = \prod_{j:s_j<t} 1 - h_j.$$

We set $c_0 = 1$.

Lastly, to estimate the posterior survival curves, we repeatedly draw from the posterior distribution of the survival probability at time t, $S^{pre}(t|x)$, for $t = 1, 2, \cdots, 120$ being each month over the 10 years.

4.2.6.2 Estimating the increased hazard due to delaying cancer treatment for pre-COVID-19 mortality

Let y be the number of days of the delay, HR_y be the hazard ratio for delaying cancer treatment by y days, and $S_i^{pre}(t|x)$ be a draw from the posterior distribution of survival pre-COVID-19 under standard treatment that has not been delayed at time t for a patient with covariates x. Then the estimated survival at time t with a treatment delay of length y is $S_{y,i}^{pre}(t|x) = S_i^{pre}(t|x)^{HR_y}$. By repeating this, we can draw from the posterior distribution of survival pre-COVID-19 for a cancer patient with delayed treatment.

We obtained estimates of the HR of delaying cancer treatment in the absence of COVID-19 from NCDB as described in Hartman et al. (2020) with collaborators from Penn State College of Medicine. We utilized the NCDB including patients treated between 2004 and 2014 (n=5,436,896).

Time to treatment from diagnosis (any treatment) was calculated for each patient. Patients were excluded if their time-to-treatment was greater than 180 days after diagnosis or missing, time to death or last contact was missing, or clinical stage was missing. Stratified Cox proportional hazard models for each cancer type (with year of diagnosis as stratum) were fit. In addition to clinical stage, the models included covariates for race (white, black, other), rurality (urban, rural), age groups (<50, 50-60, 60-70, 70-80, 80-90, >90 years old), insurance status, education, income, facility type, facility location, and distance to hospital.

4.2.7 Estimating predicted survival curves and posterior prediction intervals for overall survival under differing treatment patterns

In order to estimate the overall survival of a patient, we combined the cause-specific COVID-19 survival and the pre-COVID-19 overall survival. We assumed that failure time from COVID-19 and non-COVID-19 causes are independent conditional on the covariates we use to estimate each survival function (i.e. age, comorbidities, cancer type, and stage). Thus, the overall survival function is (Pfeiffer and Gail, 2017):

$$S(t|x) = S^{pre}(t|x)S^{COVID}(t|x).$$

We draw repeatedly from the posterior distribution of cause-specific COVID-19 survival, $S^{COVID}(t|x)$, and the posterior distribution of pre-COVID-19 overall survival, $S^{pre}(t|x)$, to simulate the posterior distribution of $S(t)$. We then selected the 2.5th percentile and the 97.5th percentile to be the posterior prediction interval.

We then calculated RMST for time for interval $[0, T]$ as:

$$\sum_{t=0}^{T} S(t|x)$$

Our primary outcome of interest is the difference in RMST between immediate (treatment beginning 7 days after diagnosis) and delayed (treatment beginning 60 days after diagnosis) cancer treatment for an individual patient based on their specific covariates.

4.2.8 Sensitivity analyses

To assess the sensitivity of our model, we conducted two sensitivity analyses. As our knowledge of the properties of COVID-19 is constantly growing, it is important that our method be robust to misspecifications. At the time of work, there were two main issues facing COVID-19 prediction

modeling. One was the estimation of the rate of vaccination. The other was the underreporting of COVID-19 infections, possibly due to asymptomatic COVID-19 cases. We explored both of these to examine the robustness of our method. We compared the difference in RMST between immediate treatment and delayed treatment under different assumptions regarding either vaccination rates or underreporting of cases.

4.2.8.1 Vaccination

Data on the number of people in a state that have received all doses of a vaccine over time were obtained from Centers for Civic Impact at Johns Hopkins (2021). Limited data is available at more granular levels. We added the most recent estimates of the proportion of currently vaccinated individuals to the recovered/removed category. We also added a transition probability from susceptible to recovered/removed as people become vaccinated. We call ξ_t the vaccination rate which can vary over time (Figure 4.4). As an approximation to the complex reality, we assumed that 75% of HCWs will be vaccinated before the general public. We fit a Bayesian linear model to number of new vaccinations administered by day and use draws from the posterior distribution for ξ_t.

4.2.8.2 Under reporting number of true COVID-19 cases

We explore the effects of under reporting the number of true infections or cases, which would impact the number of susceptible people later in the course of the pandemic. In this scenario, we start with a higher proportion of people in the recovered and infected categories. To do this, we no longer assume the mean number of observed cases (y_t) is the true number of true cases (θ^I) in the Dirichlet-Beta state space model as described in Section 4.2.2.2. Instead, we assume that $E(y_t) = \alpha\theta_I$ where $0 < \alpha < 1$ is some constant. We perform the analyses under the assumption of under reporting cases to determine the robustness of our method.

4.3 Results

4.3.1 Literature values

Since we model HCWs and the general population separately, we assume that $H = 12\%$ of the population work in healthcare (Kaiser Family Foundation, 2018). We also assume that $h = 0.23\%$ of the general population receives care in a health care setting each day (Centers for Disease Control and Prevention, 2017). We set $\beta_H = 3.47\beta_g$ (Wang, Liu, et al., 2020) and $D = 14$ (World Health Organization, 2020).

Figure 4.4: The susceptible-infected-recovered directed acyclic graph where there is a separate infection risk when at health care facilities and people become vaccinated at a set rate.

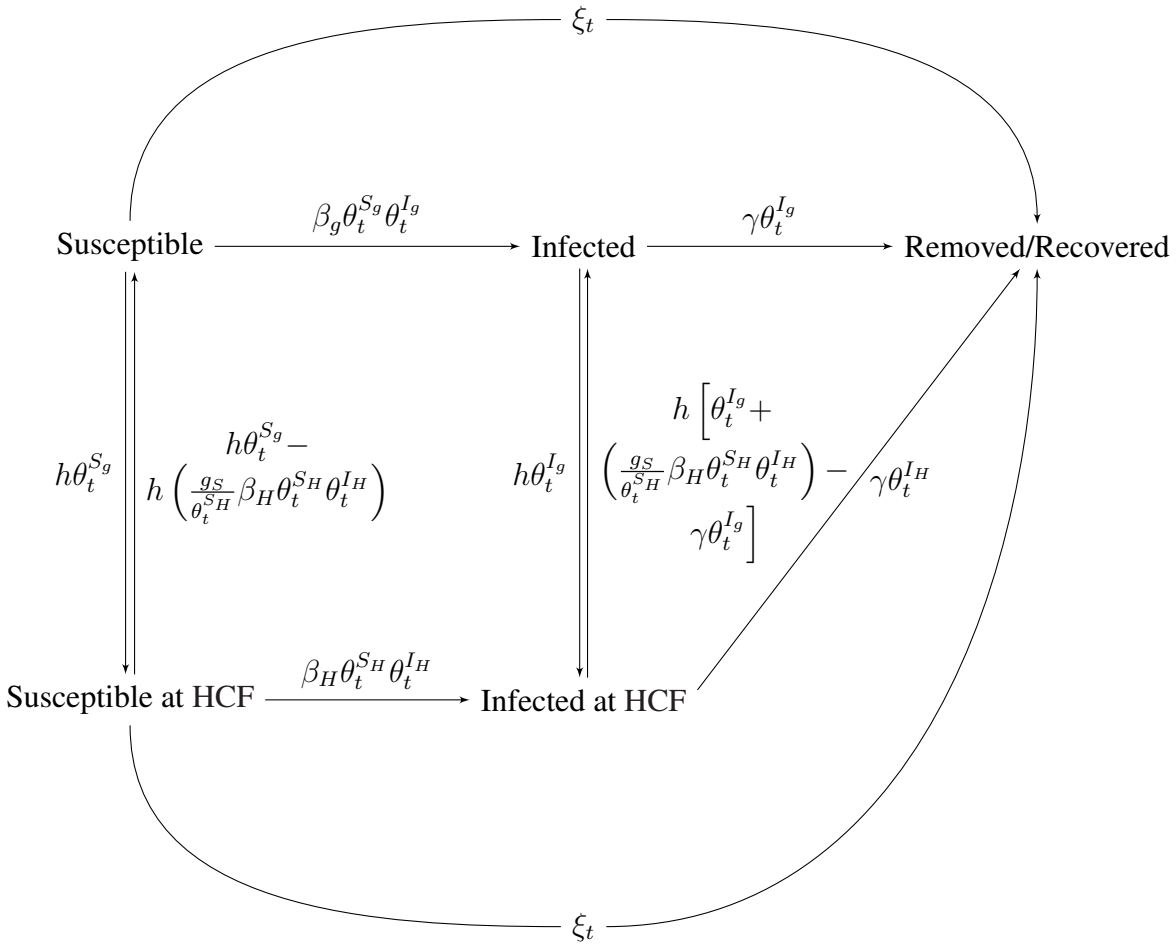

4.3.2 Validation of mpSIR model with truncated Dirichlet method

Overall, our method performed well when there was no major policy change or event that have disrupted the transmission or tracking of COVID-19 (Figure 4.5). We suggest that this model be revisited and re-run regularly when making treatment decisions.

We see that when strategies for reducing the spread of COVID-19 were implemented, our model fails to accurately predict the true number of cases (Figure 4.6). Similarly, when restrictions were lifted, the increase in cases was not accurately predicted by our model (Figure 4.6).

The IHME method aims to account for under reporting in the number of COVID-19 cases so it frequently had higher estimates than the New York Times confirmed cases and our method. The IHME method should be compared to other methods and the truth only in the shape of the curve and not necessarily the height to account for this fundamental difference in approach. Additionally, IHME only presents predictions about 110 days out, where our method makes predictions over longer time periods. However, this may not be practically beneficial if significant policy changes are expected.

4.3.3 Estimating case fatality rate of COVID-19 by number of comorbidities, age, and sex

We calculated the CFR and normalized the data. The estimates for fatality and the 95% credible intervals by age, gender, and number of comorbidities are presented in Table 4.1.

We see that there is a wide range of CFRs from 0.61% for a female between the ages of 51-60 with no comorbidities up to 70.8% for a male over 81 years old with three or more comorbidities. This wide range is consistent with other reports of COVID-19 CFRs estimated by age or by number of comorbidities (The Novel Coronavirus Pneumonia Emergency Response Epidemiology Team, 2020; Undurraga et al., 2021). The larger values for people 81+ with 3+ comorbidities are likely due to the small number of people actually in that category (Table 4.2). In our NHANES analysis, we estimated that only 2.9% of the US population that is 81 and older has 3+ comorbidities. For the age groups 81+, 71-80, 61-70, the most common number of comorbidities was 1. For the age group 51-60, the most common number of comorbidities was 0.

4.3.4 Estimating pre-COVID-19 survival with SEER data

The posterior distributions for survival by disease site and age group were calculated for the majority of cancer stages. In the case of no deaths for a given age range, site, and stage, we were unable to calculate the posterior distribution.

Two examples of the produced survival curves are shown in Figure 4.7 and Figure 4.8. There

Figure 4.5: The proportion of the general public predicted to be infected with COVID-19 as predicted by IHME and the mpSIR method. Solid line is the median and dashed lines are the 95% credible intervals. The solid blue line indicates the proportion infected as measured by confirmed cases. The plots are annotated with events that may have an impact on the number of confirmed cases. mpSIR = mixing population Susceptible-Infected-Recovered/Removed model, IHME = Institute or Health Metrics and Evaluation model, Truth = confirmed cases as reported by the New York Times.

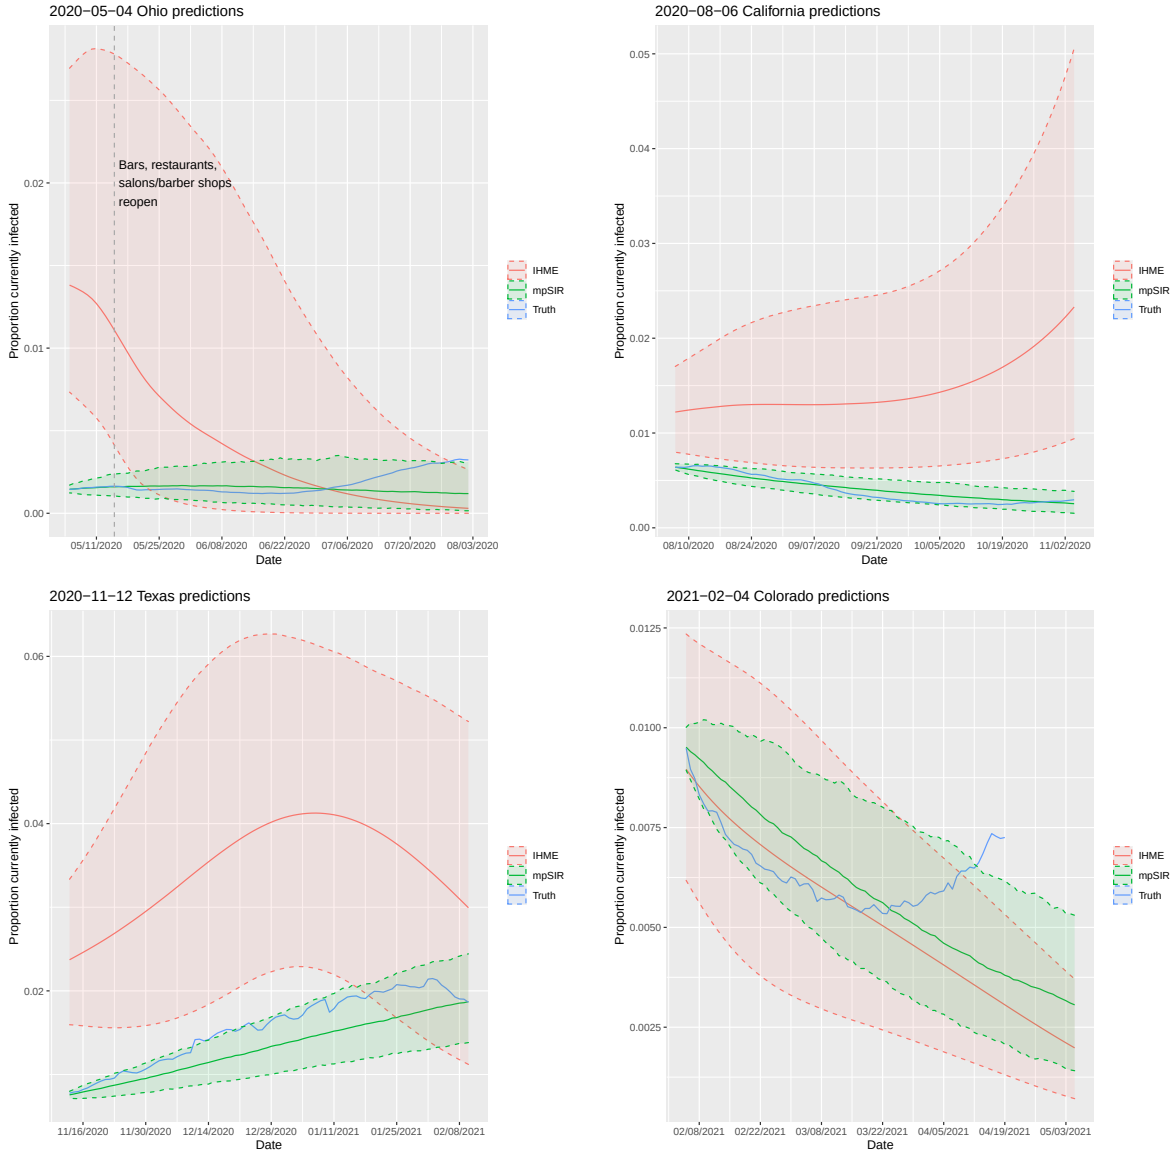

Figure 4.6: The proportion of the general public predicted to be infected with COVID-19 as predicted by IHME and the mpSIR method. Solid line is the median and dashed lines are the 95% credible intervals. The solid blue line indicates the proportion infected as measured by confirmed cases. The plots are annotated with events that may have an impact on the number of confirmed cases. mpSIR = mixing population Susceptible-Infected-Recovered/Removed model, IHME = Institute or Health Metrics and Evaluation model, Truth = confirmed cases as reported by the New York Times.

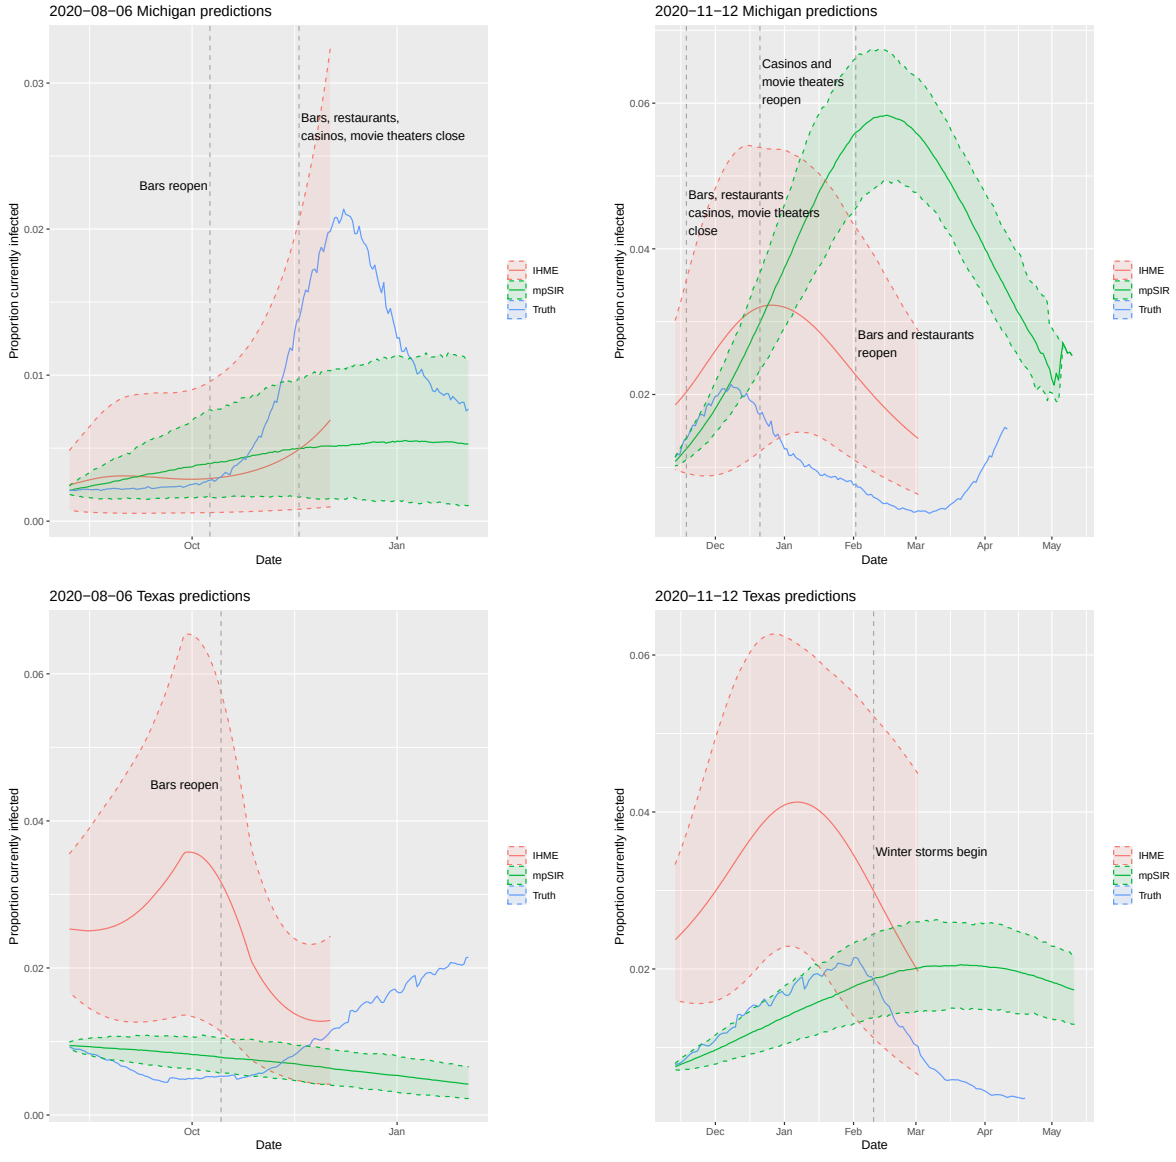

were 110,950 lung cancer patients between the ages 60 and 64 that included in the analyses that produced Figure 4.7 and there were 1,100 nasopharynx cancer patients between the ages of 60 and 64 that were included in the analyses that produced Figure 4.8. As expected, when the data used includes fewer patients with the specific type of cancer, the posterior prediction intervals are wider than when based on data with more patients.

Figure 4.7: The posterior predicted pre-COVID-19 survival curves and 95% credible interval for lung cancer patients aged 60-64 by stage.

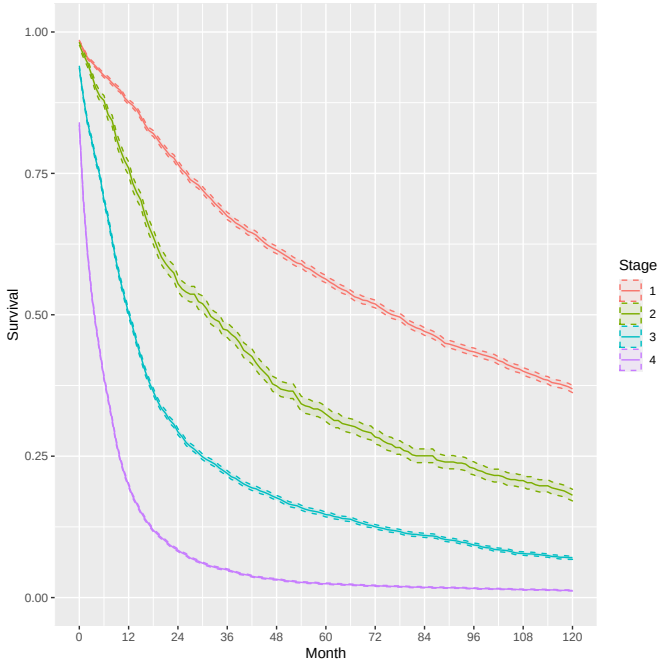

4.3.5 Combining COVID-19 survival and pre-COVID-19 overall survival estimates

We compared immediate treatment (treatment beginning 7 days after diagnosis) versus delayed treatment (treatment beginning 60 days after diagnosis). We used COVID-19 case data up until April 14, 2021 for making predictions about COVID-19 infection risk.

We used an example stage II breast cancer patient that is 64 years old with two additional comorbidities. Treatment for this patient would involve chemotherapy and surgery where she would be go into the hospital 35 times over 180 days. The hazard ratio in pre-COVID-19 survival for a 60 day delay is 1.14. The case fatality rate for a woman of age 64 with two additional comorbidities is 9.22% (95% CI: 5.79%-13.47%). Once the patient begins chemotherapy, the COVID-19 CFR increases by 1.11-fold (Lee et al., 2020). We estimated overall survival out to

Figure 4.8: The posterior predicted pre-COVID-19 survival curves and 95% credible interval for nasopharynx cancer patients aged 60-64 by stage.

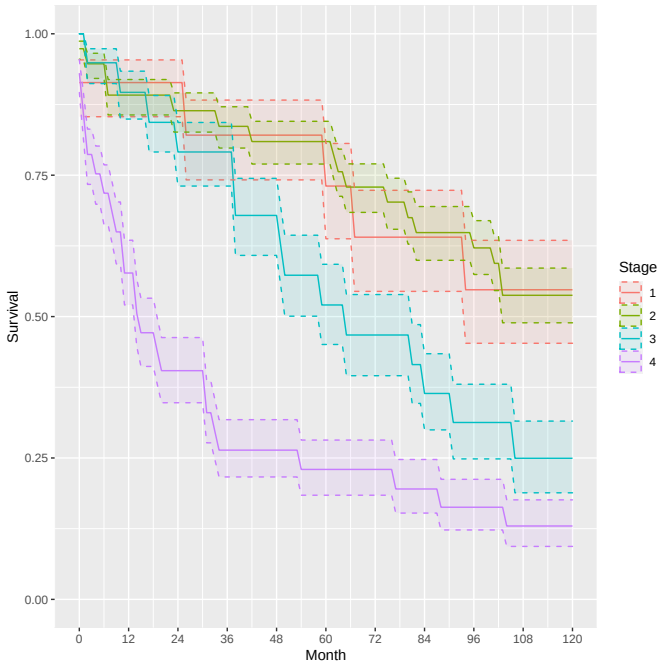

5 years for such a patient in Davidson County, TN (Nashville area) and in Denver County, CO (Denver area). These results and differences in RMST for truncated times by year up to 5-years are seen in Figure 4.9. If two patients have identical parameters and treatment plans, but are in different locations, the credible intervals change widths and the difference between treatment plans appears more stark. Since this type of patient has a high pre-COVID-19 survival rate, the effect of COVID-19 can be clearly seen in the shape of the survival curves. In Denver County, where cases were expected to increase, the patient has an immediate drop in survival probability. In Davidson County, the initial drop is not seen. Since immediate treatment is slightly preferred even when COVID-19 cases increased, this suggests that there is actually minimal risk associated with going to a HCF to receive care.

The next example patient is a stage II prostate cancer patient that is 73 years old with two additional comorbidities living in Fulton County, GA (Atlanta area). Treatment for this patient would not involve chemotherapy or surgery. The patient would go into the hospital 20 times over 28 days for radiation treatment. The hazard ratio in pre-COVID-19 survival for a 60 day delay is 1.06. Results are seen in Figure 4.9. There is very little noticeable difference between immediate and delayed treatment. There is a major risk for this patient in terms of COVID-19 mortality since they are over 70 years old with additional comorbidities and the case fatality rate is 39.31% (95% CI: 24.53%-57.88%).

Finally, we have an example male patient with stage I lung cancer that is 72 years old with two additional comorbidities also living in Fulton County, GA. Treatment for this patient would include chemotherapy and surgery and the patient would go into the hospital twice over 2 days. The hazard ratio for delaying treatment is 1.24. The survival estimates and RMSTs are seen in Figure 4.9. This patient is also at major risk for COVID-19 mortality with a COVID-19 CFR of 39.31% (95% CI: 24.53%-57.88%).

Comparing the previous two examples, we see the results of patients in the same location but with different cancer disease sites. We can see that the effect of COVID-19 is larger on the survival estimates for the patient with prostate cancer since there prognosis is generally good for stage 2 prostate cancer. However, there is a small difference in RMST at 5 years between immediate and delayed treatment (0.2 months, 95% CI: 0.19-0.21). For the lung cancer patient, there is a preference for immediate treatment and a larger difference in RMST at 5 years (3.2 months, 95% CI: 2.92-3.36). Since this patient's cancer has a lower survival rate, the effect of COVID-19 on earlier survival is not as apparent.

4.3.6 Sensitivity analyses

Overall there was very little difference in our outcome of interest, the difference in RMST under immediate versus delayed treatment, when we incorporated vaccine data (Figure 4.10) or if we assumed a level of under-reporting for COVID cases (Figure 4.11). Conversely, the predicted proportion of people with COVID-19 over time did vary substantially depending on how these factors were incorporated. As expected, when we assumed the confirmed cases was an under count of the true number of cases, we saw that the proportion infected on the first day was dramatically higher than when we assumed the confirmed cases was on average equal to the true number of cases. However, this difference did not produce a large difference in the difference in RMST. Similarly, when we took into consideration the vaccination rates of the various states, we saw that the predicted proportion of infected individuals decreased much faster than when we ignored data on vaccination rates. Again, this did not have a large impact on the difference in RMSTs by treatment timing. Overall, it seems that our proposed method is fairly robust to misspecifications in our prediction methods for COVID-19 infection risk.

The largest difference in RMST seen in these examples occurred when the true cases were assumed to be higher than the confirmed cases in Denver County, CO. We used the example breast cancer patient as described in Section 4.3.5. Assuming the confirmed cases are on average equal to the true number of cases, the difference in RMST at 5 years between immediate and delayed treatment was 0.32 months (95% CI: 0.26-0.35). When we assumed that the confirmed cases were on average 25% of the true cases, then the 5 year difference in RMST was 0.14 months (95% CI:

Figure 4.9: The posterior predicted survival curves and 95% credible interval for a lung cancer patient during COVID-19 estimated using data up until April 14, 2021. RMST is calculated with truncation time of each year up to 5 years. Solid line is the median and dashed lines are the 95% posterior prediction intervals. RMST = restricted mean survival time, yo = years old

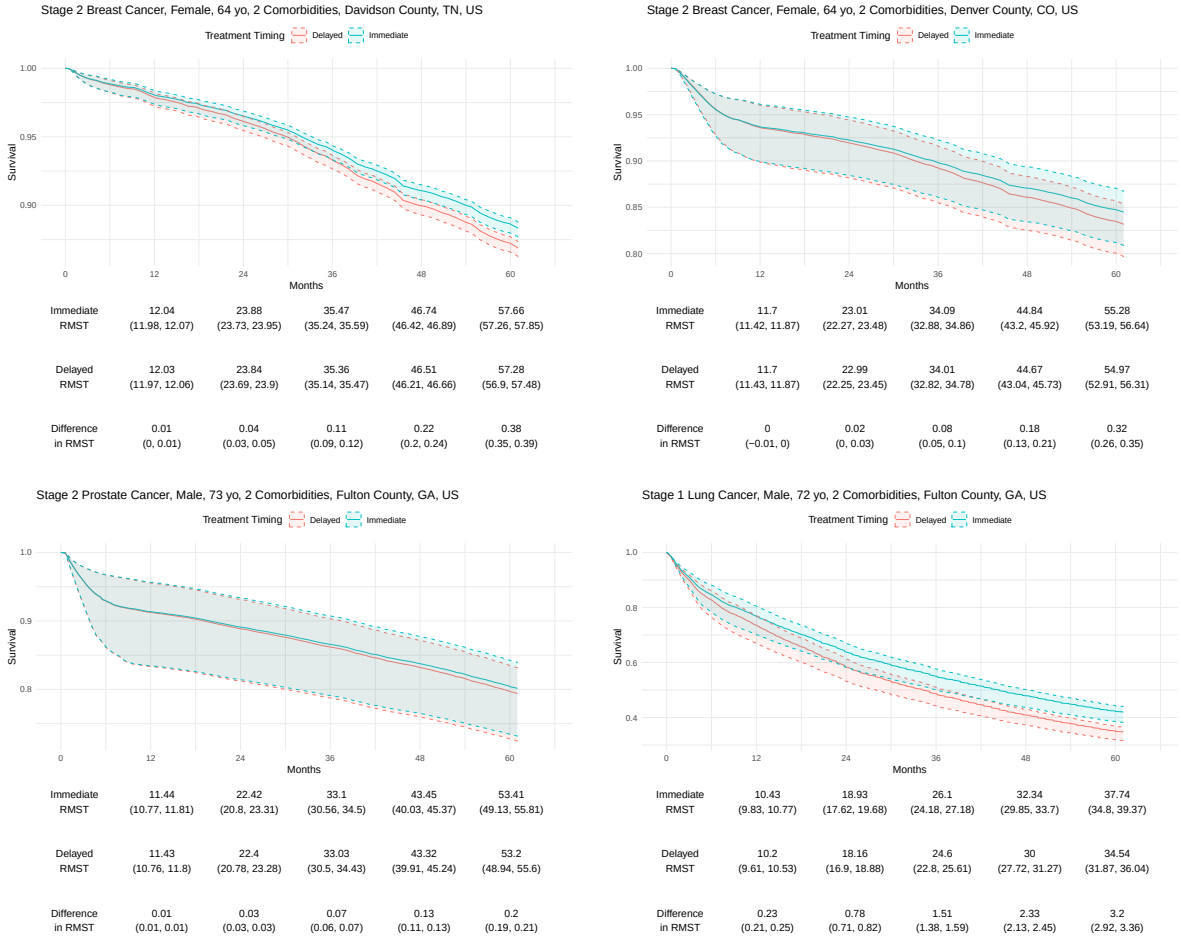

Immediate RMST	12.04 (11.98, 12.07)	23.88 (23.73, 23.95)	35.47 (35.24, 35.59)	46.74 (46.42, 46.89)	57.66 (57.26, 57.85)
Delayed RMST	12.03 (11.97, 12.06)	23.84 (23.69, 23.9)	35.36 (35.14, 35.47)	46.51 (46.21, 46.66)	57.28 (56.9, 57.48)
Difference in RMST	0.01 (0, 0.01)	0.04 (0.03, 0.05)	0.11 (0.09, 0.12)	0.22 (0.2, 0.24)	0.38 (0.35, 0.39)

Immediate RMST	11.7 (11.42, 11.87)	23.01 (22.27, 23.48)	34.09 (32.88, 34.86)	44.84 (43.2, 45.92)	55.28 (53.19, 56.64)
Delayed RMST	11.7 (11.43, 11.87)	22.99 (22.25, 23.45)	34.01 (32.82, 34.78)	44.67 (43.04, 45.73)	54.97 (52.91, 56.31)
Difference in RMST	0 (−0.01, 0)	0.02 (0, 0.03)	0.08 (0.05, 0.1)	0.18 (0.13, 0.21)	0.32 (0.26, 0.35)

Immediate RMST	11.44 (10.77, 11.81)	22.42 (20.8, 23.31)	33.1 (30.56, 34.5)	43.45 (40.03, 45.37)	53.41 (49.13, 55.81)
Delayed RMST	11.43 (10.76, 11.8)	22.4 (20.78, 23.28)	33.03 (30.5, 34.43)	43.32 (39.91, 45.24)	53.2 (48.94, 55.6)
Difference in RMST	0.01 (0.01, 0.01)	0.03 (0.03, 0.03)	0.07 (0.06, 0.07)	0.13 (0.11, 0.13)	0.2 (0.19, 0.21)

Immediate RMST	10.43 (9.83, 10.77)	18.93 (17.62, 19.68)	26.1 (24.18, 27.18)	32.34 (29.85, 33.7)	37.74 (34.8, 39.37)
Delayed RMST	10.2 (9.61, 10.53)	18.16 (16.9, 18.88)	24.6 (22.8, 25.61)	30 (27.72, 31.27)	34.54 (31.87, 36.04)
Difference in RMST	0.23 (0.21, 0.25)	0.78 (0.71, 0.82)	1.51 (1.38, 1.59)	2.33 (2.13, 2.45)	3.2 (2.92, 3.36)

Figure 4.10: Figure showing the proportion of the general public infected with COVID-19 over time. Solid line is the median and dashed lines are the 95% credible intervals. Table shows the RMST and 95% credible interval for cancer patients during COVID-19 under models that incorporate COVID-19 vaccine data and models that do not. RMST = restricted mean survival time, yo = years old

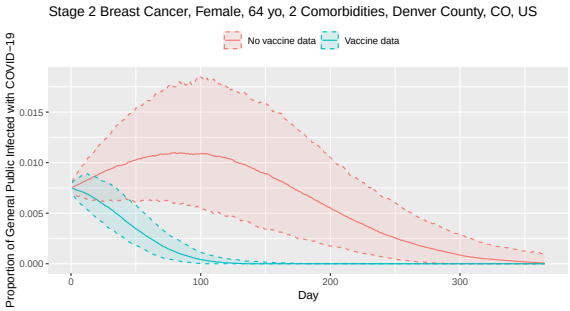

Difference in RMST for Treatment Timing

	No Vaccine Data	Vaccine Data
1 year	0 (−0.01, 0)	0.01 (0, 0.01)
2 years	0.02 (0, 0.03)	0.04 (0.03, 0.04)
3 years	0.08 (0.04, 0.1)	0.1 (0.09, 0.11)
4 years	0.18 (0.13, 0.21)	0.21 (0.2, 0.23)
5 years	0.32 (0.26, 0.35)	0.37 (0.34, 0.39)

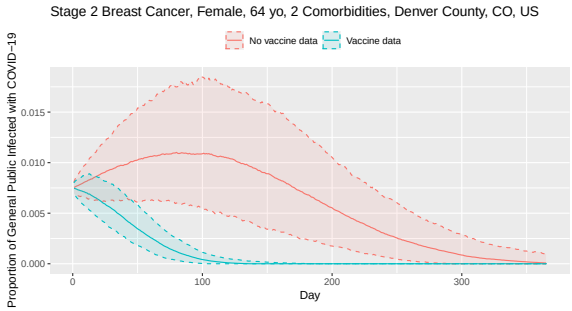

Difference in RMST for Treatment Timing

	No Vaccine Data	Vaccine Data
1 year	0 (−0.01, 0)	0.01 (0, 0.01)
2 years	0.02 (0, 0.03)	0.04 (0.03, 0.04)
3 years	0.08 (0.04, 0.1)	0.1 (0.09, 0.11)
4 years	0.18 (0.13, 0.21)	0.21 (0.2, 0.23)
5 years	0.32 (0.26, 0.35)	0.37 (0.34, 0.39)

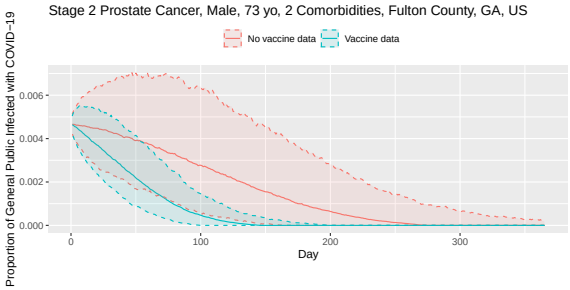

Difference in RMST for Treatment Timing

	No Vaccine Data	Vaccine Data
1 year	0.01 (0.01, 0.01)	0.24 (0.23, 0.25)
2 years	0.03 (0.03, 0.03)	0.81 (0.77, 0.84)
3 years	0.07 (0.06, 0.07)	1.57 (1.51, 1.62)
4 years	0.13 (0.11, 0.13)	2.43 (2.34, 2.5)
5 years	0.2 (0.19, 0.21)	3.34 (3.22, 3.42)

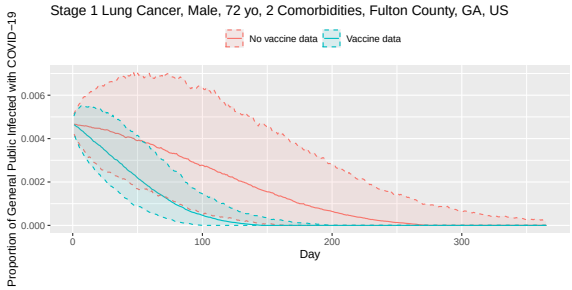

Difference in RMST for Treatment Timing

	No Vaccine Data	Vaccine Data
1 year	0.23 (0.21, 0.25)	0.24 (0.23, 0.25)
2 years	0.78 (0.71, 0.82)	0.81 (0.77, 0.84)
3 years	1.51 (1.38, 1.59)	1.57 (1.51, 1.62)
4 years	2.34 (2.13, 2.45)	2.43 (2.34, 2.5)
5 years	3.21 (2.91, 3.36)	3.34 (3.22, 3.42)

Figure 4.11: Figure showing the proportion of the general public infected with COVID-19 over time. Solid line is the median and dashed lines are the 95% credible intervals. Table shows the RMST and 95% credible interval for cancer patients during COVID-19 under models that account for under reporting COVID-19 cases. RMST = restricted mean survival time, yo = years old

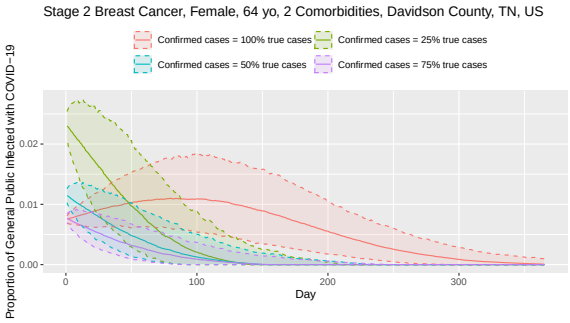

Difference in RMST for Treatment Timing

	100%	75%	50%	25%
1 year	0.01 (0, 0.01)	0.01 (0, 0.01)	0 (0, 0.01)	0 (−0.02, 0.01)
2 years	0.04 (0.03, 0.05)	0.04 (0.03, 0.05)	0.03 (0.02, 0.04)	0.02 (−0.01, 0.04)
3 years	0.11 (0.09, 0.12)	0.1 (0.09, 0.11)	0.1 (0.07, 0.11)	0.07 (0.02, 0.1)
4 years	0.22 (0.2, 0.24)	0.22 (0.19, 0.23)	0.21 (0.17, 0.23)	0.17 (0.1, 0.21)
5 years	0.38 (0.35, 0.39)	0.37 (0.34, 0.39)	0.35 (0.31, 0.38)	0.31 (0.22, 0.37)

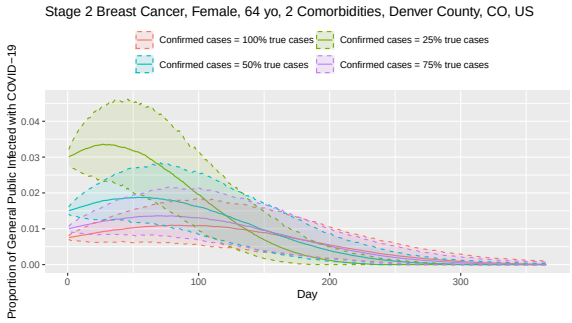

Difference in RMST for Treatment Timing

	100%	75%	50%	25%
1 year	0 (−0.01, 0)	−0.01 (−0.02, 0)	−0.01 (−0.03, 0)	−0.03 (−0.06, −0.01)
2 years	0.02 (0, 0.03)	0.01 (−0.01, 0.03)	−0.01 (−0.04, 0.02)	−0.04 (−0.1, −0.01)
3 years	0.08 (0.04, 0.1)	0.06 (0.02, 0.09)	0.03 (−0.02, 0.07)	−0.03 (−0.11, 0.03)
4 years	0.18 (0.13, 0.21)	0.16 (0.1, 0.19)	0.12 (0.04, 0.17)	0.04 (−0.08, 0.11)
5 years	0.32 (0.26, 0.35)	0.29 (0.22, 0.34)	0.24 (0.15, 0.31)	0.14 (−0.01, 0.24)

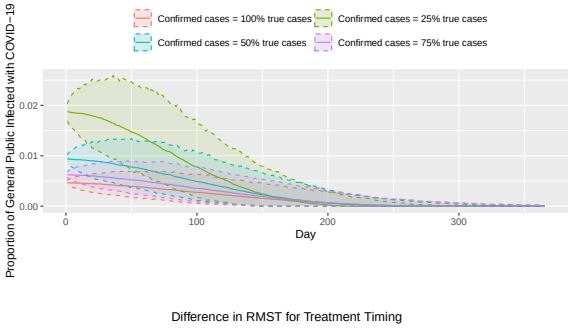

Difference in RMST for Treatment Timing

	100%	75%	50%	25%
1 year	0.01 (0.01, 0.01)	0.01 (0.01, 0.01)	0.01 (0.01, 0.01)	0.01 (0, 0.01)
2 years	0.03 (0.03, 0.03)	0.03 (0.02, 0.03)	0.03 (0.02, 0.03)	0.02 (0.02, 0.03)
3 years	0.07 (0.06, 0.07)	0.07 (0.06, 0.07)	0.06 (0.05, 0.07)	0.06 (0.04, 0.07)
4 years	0.13 (0.11, 0.13)	0.12 (0.11, 0.13)	0.12 (0.1, 0.13)	0.11 (0.08, 0.12)
5 years	0.2 (0.19, 0.21)	0.2 (0.18, 0.21)	0.19 (0.17, 0.21)	0.17 (0.13, 0.2)

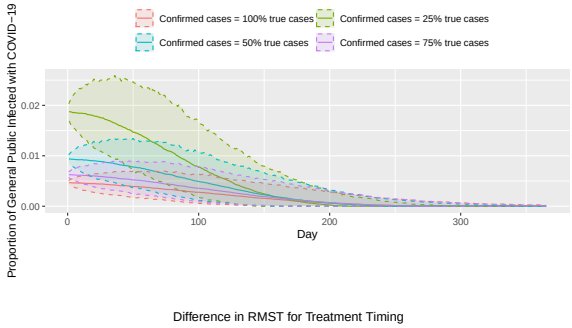

Difference in RMST for Treatment Timing

	100%	75%	50%	25%
1 year	0.23 (0.21, 0.25)	0.23 (0.2, 0.24)	0.21 (0.17, 0.23)	0.17 (0.1, 0.21)
2 years	0.78 (0.71, 0.82)	0.76 (0.67, 0.81)	0.72 (0.59, 0.79)	0.61 (0.4, 0.73)
3 years	1.51 (1.38, 1.59)	1.47 (1.3, 1.57)	1.4 (1.17, 1.53)	1.21 (0.83, 1.43)
4 years	2.34 (2.13, 2.45)	2.28 (2.03, 2.42)	2.17 (1.82, 2.37)	1.89 (1.33, 2.22)
5 years	3.21 (2.91, 3.36)	3.13 (2.79, 3.32)	2.98 (2.52, 3.26)	2.62 (1.88, 3.05)

-0.01-0.24). However, in practical terms, 0.32 months is likely not meaningfully different from 0.14 months.

4.4 Discussion

We combined Bayesian methods for several types of analyses into a new method aimed at guiding treatment decisions for clinicians and cancer patients. The proposed method draws on Bayesian methods for infectious disease prediction, survival analysis, and logistic regression. The number of analyses required for estimates for an individual patient speak to the complexity of the interactions between environment, COVID-19, and cancer. We used appropriate methods for each component to provide predictions along with their posterior prediction intervals and combined these results within a Bayesian framework.

For most patients between the options of immediate or delayed treatment, there appears to be minimal difference in RMST defined as the average survival time within a given time interval after accounting for various levels variation and uncertainty. This likely indicates that the risk of going into the HCF is not significant enough to warrant delaying cancer treatment on its own. However, if HCFs need to delay treatment due to being overwhelmed, then these methods can help prioritize patients by identifying those who's overall survival will be most impacted by delaying treatment.

While we accounted for several layers of uncertainty, there are other factors of uncertainty that we did not take into consideration. For example, we used fixed values for the HR of delaying cancer treatment, the age group CFRs to which we standardized to estimate the COVID-19 CFR, and the estimates of the proportion of the population in each age group with a given number of comorbidities. Incorporating the variability from these estimates would likely widen the credible intervals in our final estimates of overall survival and RMST. Future work could be done to incorporate uncertainty from these estimates into our method.

Another limitation to our approach is that we are not able to predict COVID-19 policy changes such as opening/closing bars or restaurants. Our method faces the same limitations as all prediction models in correctly predicting long-term risk of contracting COVID-19. Long term survival would likely need to be continually re-evaluated to take into account the latest policies and their effect on COVID-19 spread. Additionally, the literature on COVID-19 is expanding daily and although we aimed to incorporate uncertainty, there is also the possibility that new research will necessitate changes in our methods and alter these results. Despite these limitations, we show that this method is robust to misspecifications in the properties of COVID-19 detection and spread. Since there is still much to learn about the detection and spread of COVID-19, robustness is critical to making this method usable. One such example is the new Delta variant of COVID-19. This variant is causing a higher rate of reinfections than other previous variants. If reinfection becomes a more

common occurrence, then we will need to adjust our mpSIR model to allow for reinfection and remove the monotonicity constraint from our truncated Dirichlet forecasting analysis.

One difficulty in this project was due to combining multiple data sources in order to estimate overall survival during a pandemic. Cancer registries interested in being able to have relevant data for possible future pandemic or public health emergency should collect data on cancer treatment delays as well as the cause of the delay. Being able to better understand the effects of cancer treatment delay will allow more insight into the risks and benefits of delaying cancer treatment during a public health emergency. Other important data includes the cause of death. Having both the length and cause of the delay and the cause of death will allow us to see the specific effect of delaying cancer treatment on cancer specific survival. This will reduce the amount of confounding if a person's cancer treatment is delayed in order to first treat another disease which may also impact their overall survival, but not their cancer survival.

We aim to update our web app or create a secondary web app with this method to allow clinicians to easily access this model globally. The results will help guide patient treatment as the COVID-19 pandemic continues. The combination of methods presented here could provide a framework for future for integrating data in a future scenario. This may be useful in setting guidelines on managing care for high risk patients in cases of localized epidemics.

Table 4.1: Median case fatality rate (%) and 95% credible intervals by age range and comorbidities.

Male	Number of Comorbidities			
Age (years)	0	1	2	3+
51-60	1.14 (0.71, 1.9)	1.56 (1.04, 2.76)	2.14 (1.49, 3.69)	4.23 (3.04, 5.64)
61-70	4 (2.11, 6.99)	5.46 (3.05, 9.49)	7.47 (4.2, 12.97)	14.54 (9.13, 21.25)
71-80	10.95 (5.72, 19.38)	14.84 (8.35, 26.24)	20.36 (11.4, 36.14)	39.31 (24.52, 57.88)
81+	19.59 (10.84, 38.58)	26.75 (15.78, 56.34)	36.64 (21.67, 74.65)	70.75 (46.07, 98.34)
Female	Number of Comorbidities			
Age (years)	0	1	2	3+
51-60	0.61 (0.38, 1.01)	0.83 (0.56, 1.47)	1.14 (0.79, 1.96)	2.25 (1.62, 3)
61-70	2.53 (1.34, 4.43)	3.46 (1.93, 6.01)	4.73 (2.67, 8.22)	9.22 (5.79, 13.47)
71-80	7.02 (3.67, 12.43)	9.52 (5.35, 16.83)	13.06 (7.31, 23.18)	25.21 (15.72, 37.12)
81+	14.17 (7.84, 27.91)	19.35 (11.42, 40.75)	26.5 (15.67, 54)	51.18 (33.32, 71.13)

Table 4.2: Percentage of the US population in a given age bracket with 0, 1, 2, or 3+ comorbidities by age group. Rows sum to 1.

	Number of Comorbidities			
Age (years)	0	1	2	3+
40-50	0.457	0.361	0.152	0.030
51-60	0.363	0.331	0.244	0.062
61-70	0.257	0.399	0.241	0.104
71-80	0.261	0.413	0.239	0.087
81+	0.312	0.440	0.219	0.029

CHAPTER 5

Summary and Future Work

In this dissertation, I have presented a novel clinical trial design applied to small and larger sample sizes with different goals that allow greater flexibility while maintaining reliable statistical performance. These trials demonstrate how clinical trials can proceed without a dichotomous or categorical tailoring variable. Additionally, I presented a Bayesian method for combining the competing risk of an infectious disease with current data on cancer survival to develop personalized posterior survival curves with corresponding prediction intervals for cancer patients during the COVID-19 pandemic.

In Chapter 2, I described a new small sample, sequential, multiple assignment, randomized trial (snSMART) design that estimates the best first stage treatment for patient without the need for a binary outcome. The trial progresses by first randomizing patients evenly to one of three treatments. The first stage outcome is measured and a mapping function is used to assign each patient a probability of staying on the same treatment. If the patient is randomized to not stay on the first stage treatment then they are assigned to one of the remaining two treatments with equal probability. This trial leverages data from the two stages in a Bayesian framework to estimate the best first stage treatment. We show that the mapping function is more robust to misspecification than the tailoring variable. Additionally, patient outcomes on the trial were similar so adding this randomization does not result in worse patient outcomes. The mapping function used in this snSMART design has not been used in previous trials and results in a highly flexible trial design.

Motivated by the results in Chapter 2, in Chapter 3 we apply a similar concept to full sized sequential, multiple assignment, randomized trials (SMARTs) with the goal of estimating a dynamic treatment regimen (DTR). As in the snSMART trial, patients are randomly assigned to a first stage treatment with equal probability, the first stage outcome is measured, and then a tailoring function is used to assign the second stage treatment. The tailoring variable maps the first stage outcome to a probability of receiving each of the second stage treatment. The second stage treatment is then selected randomly using those probabilities. The tailoring functions presented in this chapter are a generalization of a tailoring variable which can be thought of a step-wise tailoring function. We show that tailoring functions are more flexible and robust than tailoring variables.

When applying the estimated DTR to a new population, we show that the average second stage outcomes are frequently higher than when using a tailoring variable. Only when the optimal tailoring variable is known a priori does the tailoring variable perform better than the tailoring function. In practice, it is unlikely that the true optimal cut off for the tailoring variable will be known. Thus, the tailoring function is more robust and allows researchers to estimate both the tailoring variable and develop a DTR all with in one trial framework. We also noted that Q-learning performs better than tree-based reinforcement learning (TBRL) in terms of having less bias and having higher average second stage outcomes when the estimated DTR was used in a new population of patients. As with the design in Chapter 2, this method of assigning second stage treatments has not yet been used in a SMART, but demonstrate the utility and flexibility of such a design.

The concept of using a function to map an outcome to a probability of treatment in a clinical trial has widespread implications. There are many other types of trials that require categorical variables for some aspect of the design that may benefit from keeping variables continuous. For example, basket trials assign patients to a trial arm based on biomarkers. Patients are assigned a treatment arm based on the presence or absence of particular biomarkers, and the outcome is measured. A new design for a basket trial could allow for a continuous biomarker and then use a tailoring function to assign a probability of each treatment arm. Drawing on the work presented in this dissertation in SMARTs and snSMARTs, this strategy may lead to increased efficiency. Basket trials are powered based on the interaction effects of the categorical biomarker status (positive/negative) and the treatments. Powering trials based on the interaction between a continuous biomarker and the treatments may result in a reduction in the necessary sample size.

Another opportunity to apply these concepts is in dose finding trial designs. Often in oncology, there is a continuous dose of one treatment (such as radiation) that warrants investigation rather than a set of unique treatments. While this deviates from the original concept of using a continuous variable to assign treatments, this maintains the fundamental concept of maintaining continuous variables as continuous rather than implementing a categorization strategy. Development of SMARTs and snSMARTs that allow for a continuous treatment would be immediately applicable in oncology and other fields where dose finding is of interest.

In this dissertation, I also presented expanded methods for developing personalized survival estimates for cancer patients during the COVID-19 pandemic. This work demonstrated how multiple methods and data sources can be combined to yield predictions. The overall procedure to obtain overall survival in cancer patients during the COVID-19 pandemic first involved simulating posterior distributions for the current reproduction number of COVID-19, the proportion of people infected with COVID-19, the prediction of the epidemic curve, case fatality rates (CFRs) based on multiple patient demographics, and the pre-COVID-19 predicted survival for cancer patients. This was all done for a patient's specific location, demographics, and cancer. We then simulated the

posterior distribution of COVID-19 specific survival by combining the posterior distributions for the epidemic curve, the daily probability of infection, and the relevant CFR based on the patient's demographics and treatment plans. Finally, the pre-COVID-19 cancer survival was combined with the COVID-19 specific survival to obtain a posterior distribution for survival for that specific patient.

This chapter required integrating methods from epidemiology, Bayesian survival, competing risks in survival analysis, and other fields to create cohesive predictions. These methods could be used for other infectious diseases that may pose risks to vulnerable populations that need medical care. Additionally, this chapter provided a Bayesian framework for integrating methods from different fields in order to develop predictions. Moreover, this work provided an example of estimating highly personalized survival predictions to better care for patients. The estimates from this method were based on a patient's age, sex, county, number of comorbidities, cancer type, and cancer stage.

While the clinical trial development and OncCOVID algorithm are seemingly disparate projects, the overall goal with this and all my work is to apply and develop rigorous statistical and mathematical methods to help people. By providing more and better information to patients and their clinicians, we can improve health outcomes.

Bibliography

Almirall, D., Compton, S. N., Gunlicks-Stoessel, M., Duan, N., and Murphy, S. A. (2012). Designing a Pilot Sequential Multiple Assignment Randomized Trial for Developing an Adaptive Treatment Strategy. *Statistics in Medicine*, 31 (17): 1887–1902.

Altman, D. G. (1994). Problems in Dichotomizing Continuous Variables. *American Journal of Epidemiology*, 139 (4): 442.

Altman, D. G. and Royston, P. (2006). The cost of dichotomising continuous variables. *BMJ*, 332: 1080.

American College of Surgeons (2020). *Recommendations for Prioritization, Treatment and Triage of Breast Cancer Patients During the COVID-19 Pandemic: Executive Summary*. https://www.facs.org/quality-programs/cancer/executive-summary. Accessed: 2020-04-15.

Association of Community Cancer Centers (2020). *Pandemic Planning Clinical Guideline for Patients with Cancer*. https://www.accc-cancer.org/docs/documents/cancer-program-fundamentals/oh-cco-pandemic-planning-clinical-guideline_final_2020-03-10.pdf. Accessed: 2020-04-15.

Breastcancer.org (Sept. 2014). *How Chemotherapy Affects the Immune System*. https://www.breastcancer.org/tips/immune/cancer/chemo. Accessed: 2021-02-15.

Center for Disease Control and Prevention (2020). *Interim Operational Considerations for Public Health Management of Healthcare Workers Exposed to or with Suspected or Confirmed COVID-19: non-U.S. Healthcare Settings*. https://www.cdc.gov/coronavirus/2019-ncov/hcp/non-us-settings/public-health-management-hcw-exposed.html. Accessed: 2020-07-30.

Center for Disease Control and Prevention (July 2021). *Interim Clinical Considerations for Use of COVID-19 Vaccines Currently Authorized in the United States*. hhttps://www.cdc.gov/vaccines/covid-19/clinical-considerations/covid-19-vaccines-us.html. Accessed: 20201-07-28.

Centers for Civic Impact at Johns Hopkins (Apr. 2021). *Data analysis and visualizations of daily COVID cases report*. https://github.com/govex/COVID-19.

Centers for Disease Control and Prevention (Jan. 2017). *Ambulatory Care Use and Physician office visits*. https://www.cdc.gov/nchs/fastats/physician-visits.htm. Accessed: 2020-04-15.

Centers for Disease Control and Prevention (2020). *National Health and Nutrition Examination Survey*. https://wwwn.cdc.gov/nchs/nhanes/ContinuousNhanes/Default.aspx?BeginYear=2005. Accessed: 2020-04-15.

Chakraborty, B. and Murphy, S. A. (2014). Dynamic Treatment Regimes. *Annual Review of Statistics and Its Application*, 1: 447–464.

Cori, A., Ferguson, N. M., Fraser, C., and Cauchemez, S. (2013). A New Framework and Software to Estimate Time-Varying Reproduction Numbers During Epidemics. *American Journal of Epidemiology*, 178 (9): 1505–1512.

Cox, G. F. (2018). The art and science of choosing efficacy endpoints for rare disease clinical trials. *American Journal of Medical Genetics Part A*, 176: 759–772.

Cummings, D. A. and Lessler, J. (2014). "Infectious Disease Dynamics". In: *Infectious Disease Epidemiology*. Ed. by K. E. Nelson and C. M. Williams. Burlington, MA: Jones & Bartlett Learning. Chap. 6, pp. 131–166.

Day, S., Jonker, A. H., Lau, L. P. L., Hilgers, R.-d., Irony, I., Larsson, K., Roes, K. C., and Stallard, N. (2018). Recommendations for the design of small population clinical trials. *Orphanet Journal of Rare Diseases*, 13 (195): 1–9.

Drake, C. (1993). Effects of Misspecification of the Propensity Score on Estimators of Treatment Effect. *Biometrics*, 49 (4): 1231–1236.

Fang, K.-T., Geng, Z., and Tian, G.-L. (2000). Statistical Inference for Truncated Dirichlet Distribution and Its Application in Misclassification. *Biometrical Journal*, 42 (8): 1053–1068.

Gaasterland, C. M. W., Jansen-van der Weide, M. C., du Prie-Olthof, M. J., Donk, M., Kaatee, M. M., Kaczmarek, R., Lavery, C., Leeson-Beevers, K., O'Neill, N., Timmis, O., van Nederveen, V., Vroom, E., and van der Lee, J. H. (2019). The patient's view on rare disease trial design - a qualitative study. *Orphanet Journal of Rare Diseases*, 14 (31): 1–9.

Gu, T., Chu, Q., Yu, Z., Fa, B., Li, A., Xu, L., Wu, R., and He, Y. (2020). History of coronary heart disease increases the mortality rate of COVID-19 patients: a nested case-control study. *medRxiv preprint*.

Hall, V. J., Foulkes, S., Charlett, A., Atti, A., Monk, E. J., Simmons, R., Wellington, E., Cole, M. J., Saei, A., Oguti, B., Munro, K., Wallace, S., Kirwan, P. D., Shrotri, M., Vusirikala, A., Rokadiya, S., Kall, M., Zambon, M., Ramsay, M., Brooks, T., Brown, C. S., Chand, M. A., Hopkins, S., and SIREN Study Group (2021). SARS-CoV-2 infection rates of antibody-positive compared with antibody-negative health-care workers in England: a large, multicentre, prospective cohort study (SIREN). *The Lancet*, 397 (10283): 1459–1469.

Hartman, H. E., Sun, Y., Devasia, T. P., Chase, E. C., Jairath, N. K., Dess, R. T., Jackson, W. C., Morris, E., Li, P., Hochstedler, K. A., Abbott, M. R., Kidwell, K. M., Walter, V., Wang, M., Wang, X., Zaorsky, N. G., Schipper, M. J., and Spratt, D. E. (2020). Integrated Survival Estimates for Cancer Treatment Delay Among Adults With Cancer During the COVID-19 Pandemic. *JAMA Oncology*, 6 (12).

Hartman, H. E., Tamura, R. N., Schipper, M. J., and Kidwell, K. M. (2021). Design and analysis considerations for utilizing a mapping function in a small sample, sequential, multiple assignment, randomized trials with continuous outcomes. *Statistics in Medicine*, 40 (2): 312–326.

Hilgers, R.-d., Bogdan, M., Burman, C.-f., Dette, H., Karlsson, M., König, F., Male, C., Mentre, F., Molenberghs, G., and Senn, S. (2018). Lessons learned from IDeA1: 33 recommendations from the IDeAI1-net about design and analysis of small population clinical trials. *Orphanet Journal of Rare Diseases*, 13 (77): 1–17.

Ibrahim, J. G., Chen, M.-H., and Sinha, D. (2001). Bayesian Survival Analysis. New York: Springer-Verlag New York.

Institute for Health Metrics and Evaluation (Feb. 2020). *COVID-19 Projections*. https://covid19.healthdata.org/projections. Accessed: 2021-04-12.

Kaiser Family Foundation (May 2018). *Health Care Employment as a Percent of Total Employment*. https://www.kff.org/other/state-indicator/health-care-employment-as-total/. Accessed: 2020-03-31.

Khorana, A. A., Tullio, K., Elson, P., Pennell, N. A., Grobmyer, R., Kalady, M. F., Raymond, D., Abraham, J., Klein, E. A., Walsh, R. M., Monteleone, E. E., Wei, W., Hobbs, B., and Bolwell, B. J. (2019). Time to initial cancer treatment in the United States and association with survival over time: An observational study. *PLOS One*, 14 (3): e0215108.

Kidwell, K. M., Seewald, N. J., Tran, Q., Kasari, C., and Almirall, D. (2018). Design and analysis considerations for comparing dynamic treatment regimens with binary outcomes from sequential multiple assignment randomized trials. *Journal of Applied Statistics*, 45 (9): 1628–1651.

Laber, E. B. and Zhao, Y. Q. (2015). Tree-based methods for individualized treatment regimes. *Biometrika*, 102 (3): 501–514.

Landman, A., Feetham, L., and Stuckey, D. (2020). Cancer patients in SARS-CoV-2 infection: a nationwide analysis in China. *The Lancet Oncology*, 21 (3): 335–336.

Lee, B. K., Lessler, J., and Stuart, E. a. (2009). Improved propensity score weighting using machine learning. *Statistics in Medicine*, 29 (3): 337–346.

Lee, L. Y. W., Cazier, J.-B., Angelis, V., Arnold, R., Bisht, V., Campton, N. A., Chackathayil, J., Cheng, V. W. T., Curley, H. M., Fittall, M. W., Freeman-Mills, L., Gennatas, S., Goel, A., Hartley, S., Hughes, D. J., Kerr, D., Lee, A. J. X., Lee, R. J., McGrath, S. E., Middleton, C. P., Murugaesu, N., Newsom-Davis, T., Okines, A. F. C., Olsson-Brown, A. C., Palles, C., Pan,

Y., Pettengell, R., Powles, T., Protheroe, E. A., Purshouse, K., Sharma-Oates, A., Sivakumar, S., Smith, A. J., Starkey, T., Turnbull, C. D., Várnai, C., Yousaf, N., Kerr, R., and Middleton, G. (2020). COVID-19 mortality in patients with cancer on chemotherapy or other anticancer treatments: a prospective cohort study. *The Lancet*, 395 (10241): 1919–1926.

Lei, H., Nahum-Shani, I., Lynch, K., Oslin, D., and Murphy, S. (2012). A 'Smart' Design for Building Individualized Treatment Sequences. *Annual Review of Clinical Psychology*, 8 (1): 21–48.

Liang, K.-y. and Zeger, S. L. (1986). Longitudinal Data Analysis Using Generalized Linear Models Kung-Yee. *Biometrika*, 73 (1): 13–22.

McEvoy, J. P., Lieberman, J. A., Stroup, T. S., Davis, S. M., Meltzer, H. Y., Rosenheck, R. A., Swartz, M. S., Perkins, D. O., Keefe, R. S., Davis, C. E., Severe, J., and Hsiao, J. K. (2006). Effectiveness of Clozapine Versus Olanzapine, Quetiapine, and Risperidone in Patients With Chronic Schizophrenia Who Did Not Respond to Prior Atypical Antipsychotic Treatment. *American Journal of Psychiatry*, 163 (4): 600–610.

Micheletti, R. G., Pagnoux, C., Tamura, R. N., Grayson, P. C., Mcalear, C. A., Borchin, R., Krischer, J. P., and Merkel, P. A. (2020). Protocol for a randomized multicenter study for isolated skin vasculitis (ARAMIS) comparing the efficacy of three drugs: azathioprine, colchicine, and dapsone. *Trials*, 21 (362): 1–9.

Moodie, E. E. M., Chakraborty, B., and Kramer, M. S. (2012). Q-learning for estimating optimal dynamic treatment rules from observational data. *The Canadian Journal of Statistics*, 40 (4): 629–645.

Mullins, C. D., Vandigo, J., Zheng, Z., and Wicks, P. (2014). Patient-Centeredness in the Design of Clinical Trials. *Value in Health*, 17: 471–475.

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

Murphy, S. A., Laan, M. J. van der, Robins, J. M., and CPPRG (2001). Marginal Mean Models for Dynamic Regimes. *Journal of the American Statistical Association*, 96 (456): 1410–1423.

Murray, T. A., Yuan, Y., and Thall, P. F. (2018). A Bayesian Machine Learning Approach for Optimizing Dynamic Treatment Regimes. *Journal of the American Statistical Association*, 113 (523): 1255–1267.

Nahum-Shani, I., Qian, M., Almirall, D., Pelham, W. E., Gnagy, B., Fabiano, G. A., Waxmonsky, J. G., Yu, J., and Murphy, S. A. (2012a). Experimental design and primary data analysis methods for comparing adaptive interventions. *Psychological Methods*, 17 (4): 457–477.

Nahum-Shani, I., Qian, M., Almirall, D., Pelham, W. E., Gnagy, B., Fabiano, G. A., Waxmonsky, J. G., Yu, J., and Murphy, S. A. (2012b). Q-Learning: A Data Analysis Method for Constructing Adaptive Interventions. *Psychological Methods*, 17 (4): 478–494.

New York Times (Nov. 2020). *NY Times COVID-19 Data*. https://github.com/nytimes/covid-19-data. Accessed: 2021-04-12.

Nguyen, L. H., Drew, D. A., Joshi, A. D., Guo, C.-G., Ma, W., Mehta, R. S., Sikavi, D. R., Lo, C.-H., Kwon, S., Song, M., Mucci, L. A., Stampfer, M. J., Willett, W. C., Eliassen, A. H., Hart, J. E., Chavarro, J. E., Rich-Edwards, J. W., Davies, R., Capdevila, J., Lee, K. A., Lochlainn, M. N., Varsavsky, T., Graham, M. S., Sudre, C. H., Spector, T. D., and Chan, A. T. (2020). Risk of COVID-19 among frontline healthcare workers and the general community: a prospective cohort study. *medRxiv preprint*.

Osthus, D., Kickmann, K. S., Caragea, P. C., Hidgon, D., and Del Valle, S. Y. (2017). Forecasting Seasonal Influenza with a State-Space SIR Model. *The Annals of Applied Statistics*, 11 (1): 202–224.

Pfeiffer, R. M. and Gail, M. H. (2017). Absolute Risk: Methods and Applications in Clinical Management and Public Health. Boca Raton, Florida: CRC Press LLC.

Robson, J. C., Dawson, J., Doll, H., Cronholm, P. F., Milman, N., Kellom, K., Ashdown, S., Easley, E., Gebhart, D., Lanier, G., Mills, J., Peck, J., Luqmani, R. A., Shea, J., Tomasson, G., and Merkel, P. A. (2018). Validation of the ANCA-associated vasculitis patient-reported outcomes (AAV-PRO) questionnaire. *Annals of the Rheumatic Diseases*, 77 (8): 1158–1165.

Rosenbaum, P. R. and Rubin, D. (1983). The central role of the porpensity score in observational studies for causal effects. *Biometrika*, 70 (1): 41–55.

Rosenberger, W. F. (1993). Asymptotic Inference with Response-Adaptive Treatment Allocation Designs. *The Annals of Statistics*, 21 (4): 2098–2107.

Schellingerhout, J. M., Heymans, M. W., Vet, H. C. W. D., Koes, B. W., and Verhagen, A. P. (2009). Categorizing continuous variables resulted in different predictors in a prognostic model for nonspecific neck pain. *Journal of Clinical Epidemiology*, 62 (8): 868–874.

Senn, S. (2016). Mastering variation: variance components and personalised medicine. *Statistics in Medicine*, 35 (7): 966–977.

Sinha, D. and Dey, D. K. (1997). Semiparametric Bayesian Analysis of Survival Data. *Journal of the American Statistical Association*, 92 (439): 1195–1212.

Snapinn, S. M. and Jiang, Q. (Oct. 2007). Responder analyses and the assessment of a clinically relevant treatment effect. *Trials*, 8 (31).

Tamura, R. N., Huang, X., and Boos, D. D. (2011). Estimation of treatment effect for the sequential parallel design. *Statistics in Medicine*, 30 (30): 3496–3506.

Tamura, R. N., Krischer, J. P., Pagnoux, C., Micheletti, R., Grayson, P. C., Chen, Y.-F., and Merkel, P. A. (2016). A Small n Sequential Multiple Assignment Randomized Trial Design For Use in Rare Disease Research. *Contemporary Clinical Trials*, 46: 48–51.

Tao, Y., Wang, L., and Almirall, D. (2018). Tree-Based Reinforcement Learning for Estimating Optimal Dynamic Treatment Regimes. *Annals of Applied Statistics*, 12 (3): 1914–1938.

Thall, P. F., Logothetis, C., Pagliaro, L. C., Wen, S., Brown, M. A., Williams, D., and Millikan, R. E. (2007). Adaptive therapy for androgen-independent prostate cancer: a randomized selection trial of four regimens. *J Natl Cancer Inst*, 99 (21): 1613–22.

Thall, P. F. (2015). "SMART Design, Conduct, and Analysis in Oncology". In: *Adaptive Treatment Strategies in Practice*. Ed. by M. R. Kosorok and E. E. M. Moodie. SIAM. Chap. 4, pp. 41–54.

The Novel Coronavirus Pneumonia Emergency Response Epidemiology Team (2020). The Epidemiological Characteristics of an Outbreak of 2019 Novel Coronavirus Diseases (COVID-19) - China, 2020. *China CDC Weekly*, 2: 1–10.

Tudur Smith, C., Williamson, P. R., and Beresford, M. W. (2014). Best Practice & Research Clinical Rheumatology Methodology of clinical trials for rare diseases. *Best Practice & Research Clinical Rheumatology*, 28: 247–262.

Undurraga, E. A., Chowell, G., and Mizumoto, K. (2021). COVID-19 case fatality risk by age and gender in a high testing setting in Latin America: Chile, March-August 2020. *Infect Dis Poverty*, 10 (1): 11.

US Census (2020). *US Census 2019 County Estimates*. https://www2.census.gov/programs-surveys/popest/datasets/2010-2019/counties/totals/. Accessed: 2020-04-15.

US Department of Health and Human Services Food and Drug Administration (2009). *Guidance for Industry: Patient Reported Outcomes: Use in Medical Product Development to Support Labeling Claims*. Tech. rep. December. Food and Drug Administration.

Wang, C., Liu, L., Hao, X., Guo, H., Wang, Q., Huang, J., He, N., Yu, H., Lin, X., Pan, A., Wei, S., and Wu, T. (2020). Evolving Epidemiology and Impact of Non-pharmaceutical Interventions on the Outbreak of Coronavirus Disease 2019 in Wuhan, China. *medRxiv preprint*.

Wang, L., Zhou, Y., He, J., Zhu, B., Wang, F., Tang, L., Eisenberg, M., and Song, P. X. K. (2020). An epidemiological forecast model and software assessing interventions on COVID-19 epidemic in China. *medRxiv preprint*.

Wei, B., Braun, T. M., Tamura, R. N., and Kidwell, K. M. (2020). Sample Size Determination for Bayesian Analysis of small n Sequential, Multiple Assignment, Randomized Trials (snSMARTs) with Three Agents. *Journal of Biopharmaceutical Statistics*, 30 (6): 1109–1120.

Wei, B., Braun, T. M., Tamura, R. N., and Kidwell, K. M. (2018). A Bayesian analysis of small n sequential multiple assignment randomized trials (snSMARTs). *Statistics in Medicine*, 37 (26): 3723–3732.

Wei, J. T., Dunn, R. L., Litwin, M. S., Sandler, H. M., and Sanda, M. G. (2000). Development and validation of the Expanded Prostate Cancer Index Composite (EPIC) for comprehensive

assessment of health-related quality of life in men with prostate cancer. *Urology*, 56 (6): 899–905.

Williams, M., Le Calvez, K., Mi, E., Chen, J., Dadhania, S., and Pakzad-Shahabi, L. (2020). Estimating the Risks from COVID-19 Infection in Adult Chemotherapy Patients. *medRxiv preprint*.

World Health Organization (Mar. 2020). *Considerations for quarantine of individuals in the context of containment for coronavirus disease (COVID-19)*. Tech. rep. World Health Organization.

Wu, F., Laber, E. B., Lipkovich, I. A., and Severus, E. (2015). Who will benefit from antidepressants in the acute treatment of bipolar depression? A reanalysis of the STEP-BD study by Sachs et al. 2007, using Q-learning. *International Journal of Bipolar Disorders*, 3 (1): 7.

Zajonc, T. (2012). Bayesian Inference for Dynamic Treatment Regimes: Mobility, Equity, and Efficiency in Student Tracking. *Journal of American Statistical Association*, 107 (497): 80–92.

Zhao, Y., Zeng, D., Rush, A. J., and Kosorok, M. R. (2012). Estimating Individualized Treatment Rules Using Outcome Weighted Learning. *Journal of the American Statistical Association*, 107 (449): 1106–1118.

Zhou, F., Yu, T., Du, R., Fan, G., Liu, Y., Liu, Z., Xiang, J., Wang, Y., Song, B., Gu, X., Guan, L., and Wei, Y. (2020). Clinical course and risk factors for mortality of adult inpatients with COVID-19 in Wuhan, China: a retrospective cohort study. *The Lancet*, 395 (10229): 1054–1062.

Zhou, X., Wang, Y., and Zeng, D. (2018). Outcome-Weighted Learning for Personalized Medicine with Multiple Treatment Options. *Proc Int Conf Data Sci Adv Anal*, 2018: 565–574.