INCORPORATING PARTIAL ADHERENCE INTO THE PRINCIPAL
STRATIFICATION ANALYSIS FRAMEWORK

by

ERIC SANDERS

B.Sc. Queen's University, 2017

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF

MASTER OF SCIENCE

in

THE FACULTY OF GRADUATE AND POST-DOCTORAL STUDIES
(Statistics)

THE UNIVERSITY OF BRITISH COLUMBIA
(Vancouver)

August 2019

©Eric Sanders, 2019

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the thesis entitled:

Incorporating partial adherence into the principal stratification analysis framework

Submitted by Eric Sanders in partial fulfillment of the requirements for

the degree of Master of Science

in Statistics

Examining Committee:

Co-supervisor: Paul Gustafson, Statistics

Co-supervisor: M. Ehsan Karim, Population and Public Health

Abstract

Participants in pragmatic clinical trials often partially adhere to treatment. In the presence of partial adherence, simple statistical analyses of binary adherence (receiving either full or no treatment) introduce biases. We developed a framework which expands the principal stratification approach to allow partial adherers to have their own principal stratum and treatment level. We derived consistent estimates for bounds on population values of interest. A Monte Carlo posterior sampling method was derived that is computationally faster than Markov Chain Monte Carlo sampling, with confirmed equivalent results. Simulations indicate that the two methods agree with each other and are superior in most cases to the biased estimators created through standard principal stratification. The results suggest that these new methods may lead to increased accuracy of inference in settings where study participants only partially adhere to assigned treatment.

Lay Summary

Clinical trials attempt to identify when certain treatments cause outcomes that are better on average than some alternative. There are a lot of problems that can pop up when doing clinical trials, which might cause results to be inaccurate. The key goal in this thesis is to offer clinical researchers a new method to analyze clinical trial data without bias when faced with the problem of patients only partially following a prescribed treatment, particularly in the context of when the clinical trial is being done in a ‘real world’ setting where researchers might not be able to meticulously gather a great deal of patient information.

Preface

This dissertation is original, unpublished, independent work by the author, Eric Sanders.

Original motivation for the problem of addressing partial adherence in pragmatic clinical trials arose from the BC Support Unit Real-World Clinical Trials Methods Cluster, Project #2, headed by Ehsan Karim.

Contents

Abstract	iii
Lay Summary	iv
Preface	v
Table of Contents	vi
Contents	vi
List of Tables	viii
List of Figures	ix
Acknowledgements	x
1 Introduction	1
1.1 Patient-Oriented Research	1
1.2 Pragmatic Clinical Trials	1
1.3 Standard Clinical Trial Analytical Methods	3
1.3.1 Intent-to-Treat Analysis	4
1.3.2 As-Treated Analysis	5
1.3.3 Per-Protocol Analysis	6
1.4 Principal Stratification Analysis in the Pragmatic Clinical Trial Setting	7
1.5 Problems Introduced by Partial Adherence to Assigned Treatment	8
2 Incorporating Partial Adherence into Principal Stratification Analysis	10
2.1 Framework Setup for Partial Adherence in Principal Stratification Analysis	10
2.2 Possible Assumptions and Their Implications	14
2.2.1 Bounds Resulting From the Monotone-within-Treatment Assumption	15

2.3	A Monte Carlo Method of Generating Posterior Samples of Effect Estimates	18
3	Examples and Simulation Results	24
3.1	Limiting Distributions of Parameter Estimates	27
3.1.1	Simulation 1: Both Monotonicity Types	28
3.1.2	Simulation 2: Neither Monotonicity Type	32
3.2	Validation Trial of the Monte Carlo Method through Comparison with JAGS	36
3.3	Simulation Set Comparing to the Case Discussed by Shrier	40
4	Conclusions, Future Work	47
	References	51

List of Tables

1	Table of parameters describing data structure.	14
2	Mean square error of various inference methods and across sample sizes of 1000 and 8000.	47

List of Figures

1	Posterior sample and limiting distribution for partial treatment effect on partial compliers, monotonicity types 1 & 2.	29
2	Posterior sample and limiting distribution for full treatment effect on full compliers, monotonicity types 1 & 2.	30
3	Posterior sample and limiting distribution for full treatment effect on full population, monotonicity types 1 & 2.	31
4	Posterior sample and limiting distribution for partial treatment effect on partial compliers, no monotonicity types.	33
5	Posterior sample and limiting distribution for full treatment effect on full compliers, no monotonicity types.	34
6	Posterior sample and limiting distribution for full treatment effect on full population, no monotonicity types.	35
7	JAGS Validation, posterior sample for partial treatment effect on partial compliers, monotonicity type 2.	38
8	JAGS Validation, posterior sample for full treatment effect on full compliers, monotonicity type 2.	38
9	JAGS Validation, posterior sample for full treatment effect on full population, monotonicity type 2.	39
10	Boxplots of treatment effect point estimates over 2000 simulations with sample size 1000.	43
11	Bar plot of treatment effect 95% interval coverage over 2000 simulations with sample size 1000.	44
12	Boxplots of treatment effect point estimates over 2000 simulations with sample size 8000.	45
13	Bar plot of treatment effect 95% interval coverage over 2000 simulations with sample size 8000.	46

Acknowledgements

I thank the faculty, staff, and students of the UBC Department of Statistics, who have taught me a great deal during my years of study. Their motivating support helped in the production of this thesis, and the breadth of knowledge within the department knows no bounds. Special thanks to Paul Gustafson for patiently teaching me so many new concepts, and to Ehsan Karim for always having constructive and helpful feedback on the work I had completed.

I also offer gratitude to my parents and siblings for their unending moral support over the years of my education while I have been away from home.

1 Introduction

1.1 Patient-Oriented Research

There are records of humans using experiments to inform medical and public health decisions for more than two thousand years [4]. In more recent history, due to scientific advancement, there can often be discrepancies between the understanding and goals of clinical professionals and those of care receivers. While innovatively complex research can lead to large advancements in the fields of medicine and public health, there have been calls for support of more ‘patient-oriented research’ by organizations such as the US Patient-Centered Outcomes Research Institute, the Canadian Institutes of Health Research, and the UK National Institute for Health Research [38, 40, 14].

The precise definition of ‘patient-oriented research’ may differ between organizations, but some tenets are consistent, such that research primarily focuses on supplying whatever information is directly relevant to and desired by patients who are potentially going to receive a treatment. Furthermore, patient-oriented research aims to form observations and conclusions that incorporate such factors as patient burdens, service availability, and other diverse factors that make up an entire treatment plan [38, 40, 14]. By fulfilling these goals, patients and their caregivers can be informed to the same degree and make the best choices available to them without any confusion. Many factors are relevant when discussing whether research is ‘patient-oriented,’ but some distinctions exist in study design that can make this simpler. Pragmatic clinical trials, as the relevant example to be further discussed, may be viewed as having some patient-oriented characteristics.

1.2 Pragmatic Clinical Trials

In a standard randomized clinical trial (RCT), it is most common to have highly experienced investigators working with highly selected participants, following strict rules with detailed

follow-up schedules in order to determine the efficacy of a new intervention. A concern that may arise from a standard RCT such as this is that the results may not be directly applicable in regular clinical practice. This is because study designs tend to promote internal validity (ability to accurately inform on trends in the rigidly defined trial population) more than external validity (ability to accurately inform on trends in the total population). Pragmatic trials, originally described by Schwartz and Lellouch in 1967[45], are trials which loosen some of the strict requirements of a standard RCT in an attempt to solve this problem and promote results that are more directly applicable to a larger and more realistic population. Standard RCTs may not identify highly relevant details, such as if less experienced medical staff are likely to make mistakes administering treatment, or if patients outside strict inclusion criteria can have adverse effects, or if patients in a general clinical setting are likely to inadequately follow guidelines for treatment. Pragmatic trials thus aim to overcome this by being more realistic than standard RCTs from the very start in terms of the patients who come for treatment, the medical professionals offering treatment, and the steps necessary to receive treatment. Looking at it this way, it is clear that pragmatic clinical trials are a step towards more patient-oriented research.

Trials that place treatment and assessment responsibility on everyday caregivers, or are less restrictive and more holistic when it comes to trial participation and selection, or require less intensive and demanding participation in study proceedings and data collection are often referred to as pragmatic [11]. While the distinction tends to be loose between a standard RCT and a pragmatic trial, the PRECIS (Pragmatic-Explanatory Continuum Indicator Summary) tool has been created in an attempt to measure the level of pragmatism of a trial [49]. While pragmatism in the clinical trial setting was first described over 50 years ago, only in recent years has it been gaining popularity [41], and there are multiple studies of note in recent years that have pragmatic features [23, 44, 13, 8, 43, 6, 46]. One of the main reasons that pragmatic trials are gaining popularity is that they are better suited than

RCTs to answer questions about the feasibility or benefit of fully employing an intervention in the most ‘everyday’ clinical setting.

By making a study more pragmatic in nature, it is a much smaller step to infer what would happen if a new intervention were implemented on a large scale, and how this intervention may best benefit the population. As will be seen, this comes with important considerations.

1.3 Standard Clinical Trial Analytical Methods

Whether a clinical trial is considered pragmatic or not, there are a number of options for how to analyze collected data. Primarily, the question arises of how to handle situations where the randomization of intervention assignment to patients is threatened. The major concern to be addressed moving forward is how to proceed with useful analysis when patients do not conform to prescribed interventions, and instead behave in a way that more closely fits the description of a different intervention in the study, or none at all. As an example, a trial comparing compound Salvia pellets to nitrates as treatments for chronic angina pectoris may begin by randomizing intervention assignment to patients [28], but after data collection, it might become apparent that some patients assigned compound Salvia pellets, in fact, had access to and made some use of nitrates. This type of situation is very common and of a considerable concern because different approaches in analysis can lead to changes in conclusions. Moving forward, some of the classic approaches to a problem such as this are summarized, as well as notable benefits and drawbacks to applying each method. In many applications, it is currently preferred to present results of multiple of the below analysis types together. Each analytical approach has its weaknesses, so a certain strength is gained when different analyses agree in their findings.

1.3.1 Intent-to-Treat Analysis

In an intent-to-treat analysis, data are analyzed as if patients always received the intervention they were assigned. In other words, no changes are made in the analysis if it is revealed that a patient assigned intervention A actually received intervention B, and the data would record that the patient received intervention A. In other words assigned intervention groups are considered immutable post-randomization, regardless of intervention received[19]. Utilizing this approach, the most notable result is that the randomization process of assigning treatments is not violated, and treatment assignment can still be considered truly random, granted that the study was designed correctly. However, as a result of some people in each group potentially receiving the intervention of a different group, results tend to bias towards the null, indicating a smaller or less significant difference between intervention outcomes (though some cases exist where bias may be away from the null [48]). Here the use of the term ‘bias’ is applicable only depending on perspective. If one is interested in analyzing ‘the effect of being assigned an intervention compared to a control,’ there is no bias towards the null in using the intent-to-treat analysis, but if one is interested in analyzing ‘the effect of receiving an intervention compared to a control,’ then the intent-to-treat analysis could possibly be biased, likely towards there being no effect [48]. These two estimable quantities are in fact quite distinct and should not be confused with each other, because the first of the above effects is heavily impacted by intervention program implementations, while the second is an inherent property of a new intervention. Both of these can be highly relevant, and care is needed to tease the two apart and acknowledge that the intent-to-treat analysis can not perfectly capture the effect of actually receiving treatment.

As an example of when intent-to-treat analysis has a clear flaw, suppose that a study was poorly designed and did not adequately give participants an incentive to stick to an intervention. If the control group is doing nothing, but the treatment group has an intensive intervention plan without adequate incentive supplied, many people in the treatment group

may cease treatment, to a point where a significant proportion of people in the ‘treatment’ arm of the study are doing nothing, just like in the control arm. As a result, the analysis may indicate that the treatment is ineffective, even if the treatment would have been highly effective if enough incentive was given for people to adhere. This type of problem is even more common in the pragmatic trial context, because there are many factors that are free to impact whether a participant will carry through with proper assigned treatment. As a result, using intention-to-treat analysis in a pragmatic trial can seriously threaten the accuracy of results, because a higher proportion of patients can be expected not to follow treatment regimens perfectly.

Acknowledging that the intention-to-treat approach can accurately capture the effect of being assigned an intervention, but becomes biased when estimating the effect of receiving an intervention, it is important to note what problems may arise as soon as analysis departs from the intent-to-treat approach. As soon as one decides to alter a data set based on adherence (such as in as-treated analysis or per-protocol analysis discussed below), the randomization process has been interfered with, and as a result there is an opportunity for notable differences to arise, either between the control group and intervention group, or in the definition of the target population [27]. These can lend significant bias to estimates if not controlled for.

1.3.2 As-Treated Analysis

An as-treated analysis, as the name suggests, considers all participants to belong to the intervention group that most closely matches the treatment they actually received, regardless of what treatment they were assigned. If the two final treatment groups were appropriately similar with respect to relevant pre-treatment factors, the results could accurately estimate the effect of receiving treatment. However, the issue is that as-treated analysis is highly likely to introduce distinct differences between treatment groups, because there are factors

that impact adherence, and thus impact which intervention group patients belong to. Some of these factors are often thought to be immeasurable, such as personality traits, but it is also quite possible for the intervention itself to affect adherence. Since the two patient groups were not defined through a randomization process, there is a high risk of bias that is difficult to quantify.

As an example of when a flaw might arise in as-treated analysis, suppose a large fraction of people receiving an intervention had a negative side effect, and many of these people may cease treatment. As a result of these people being treated as members of the control group, only people who responded well to the intervention are left in the treatment group, which will bias results towards having a positive outlook on the treatment.

1.3.3 Per-Protocol Analysis

In 1980, an article from the Coronary Drug Project reported that when assigned a placebo pill, adherers to the protocol had a lower 5-year mortality risk than non-adherers, which frightened many researchers away from performing analyses that separate adherers from non-adherers [9]. In a per-protocol analysis, participants are only retained in the data set for as long as they adhere to the intervention they were assigned. Thus, someone who never adheres to their assigned intervention will be completely removed from the study for all analyses. After the mentioned 1980 article was published, per-protocol analyses were regarded very critically, since the approach separates the study population from the target population according to whatever factors (unmeasured and measured) determine adherence, some of which may be related to the effect of the intervention [25, 26]. However, some methods have been developed to update per-protocol effect estimates by adjusting for pre-randomization and post-randomization variables thought to impact adherence [37, 26], with notable improvement in results when revisiting previous studies (including the results from the 1980 article regarding the Coronary Drug Project). The development of these techniques

has led to suggestions that adjusted per-protocol analyses be presented alongside intent-to-treat analyses, so as to compare valid estimates of the effect of receiving treatment to the effect of being assigned treatment [24].

1.4 Principal Stratification Analysis in the Pragmatic Clinical Trial Setting

The three standard approaches above are the most straightforward methods available for analyzing clinical trial data, and it is seen that each has drawbacks and limitations. While the intent-to-treat approach is most popular, it is often combined with the per-protocol approach in order to actually comment on the effect of receiving treatment, as opposed to simply being assigned treatment. However, more complicated methods do exist to estimate treatment effects in clinical trials. Some of the more sophisticated methods include propensity scoring, instrumental variable analysis, and principal stratification [47, 15, 2, 22, 29]. Of interest here is the latter method, principal stratification, which will be extended in this thesis to combat the issue of partial adherence to assigned treatment, which is currently not addressed by any of the mentioned analytical methods.

Principal stratification relies on characterizing participants not based on a selection of baseline features, but based on how one variable is determined by another, and by referring to each group of similarly behaving individuals as a ‘principal stratum’ [42]. For example, if one has two binary variables, one for treatment assigned and one for treatment received, then there are four possible functions from one to the other (always being equal to the other, always equalling one of the two values, or always being the opposite value as the other). It is supposed that each participant follows one of these four functions, or in other words, belongs to one of these four principal strata. This can be put in more plain language in this example by supposing that one treatment arm can be called the ‘active intervention,’ while the other is the ‘control intervention,’ and then saying that participants either always comply, always

defy, always take the active intervention, or always take the control intervention. After supposing that these strata exist, it is possible to form estimates of treatment effects within each principal stratum, and to form estimates of the distribution of the population across principal strata. This type of analysis is commonly used in trials with adherence issues in binary treatment settings, and has been used to form effect estimates for over twenty years [1, 30, 12, 16, 3, 7, 35].

The principal stratification approach is also particularly enticing in the pragmatic trial setting because of the more nuanced inferences that can be gained from the analysis. An overall estimate can be formed of the effect of treating a whole population with the active intervention compared to the control intervention, but individual estimates can also be formed per principal stratum, such as the estimate of treatment effect among those who comply with treatment assignments. These are arguably more informative analyses than simple intent-to-treat or per-protocol approaches. Often the estimates for a principal stratification analysis are calculated through simple arithmetic based on the groups that can be clearly identified, but more in-depth estimates can be formed through Bayesian inference to give explicit interval estimates and posterior densities [47, 29].

1.5 Problems Introduced by Partial Adherence to Assigned Treatment

All previously discussed methods are faced with obstacles when clinical trial participants only partially adhere to an assigned intervention, placing them somewhere between receiving no treatment and receiving full treatment. Since all of the general methods for analyzing clinical trial data assume some dichotomous (or at least discrete) indicator for treatment received, decisions must be made on what to do with the data in order to make it suitable for the chosen analytical method. In most applications, the question has become whether to count partial compliance as full compliance or non-compliance, and efforts have been lacking

in fully accounting for the fact that some participants satisfied some (but not all) specified components of the prescribed intervention [47, 10, 50].

The most naïve approach to handle the problem of partial adherence is to simply treat partially adherent participants as either fully compliant or non-compliant. Shrier et. al. [47] assessed the effects of analyses which designate a partially compliant person as fully non-compliant or fully compliant. It was found that treating partially compliant participants as fully compliant and proceeding with principal stratification analysis led to the same kinds of biases as intent-to-treat analysis (generally towards the null); on the other hand, treating partially compliant participants as non-compliant and proceeding with principal stratification analysis led to bias away from the null when estimating treatment effect in fully compliant participants [47]. These results indicate a need for some other way to properly analyze clinical trial data when some participants only partially adhere. Jin and Rubin [31] reported a method of principal stratification which allowed for compliance to be measured continuously, where a participant’s principal stratum was defined using four continuous variables measuring the proportion of each of the two treatments received in each of the two treatment assignment settings. The method of Jin and Rubin relies on multiple steps of Bayesian sampling to estimate dose-response effects, and has since been used in multiple studies where exact proportions of treatments received may be quantified [17, 32, 39], some of which extend the methods to incorporate copula analysis [36, 5]. With the potential errors introduced by not recognizing partial compliance, and the rigorous information keeping necessary to perform the analytical methods of Jin and Rubin, there is no method of analysis for the case of partial compliance in highly pragmatic settings where detailed compliance information may not be available.

Guo [18] presented a framework that used a trichotomous variable for treatment received that included a level for ‘partial’ treatment as well as full treatment and no treatment. Propensity

scoring methods are used to form a variety of estimates regarding different type of defined compliers. In this thesis, I define a framework that also uses a trichotomous variable for treatment received, and I present simple bound estimators which consistently estimate true bounds on treatment effects in relevant population groups. These bounds could have applications as simpler inference tools that do not require a complex Bayesian formulation. I also present a Bayesian inference method of forming posterior samples of treatment effects in relevant population groups, which uses a derived sampling algorithm that is less complex than the methods of Jin and Ruben as well as Guo.

2 Incorporating Partial Adherence into Principal Stratification Analysis

2.1 Framework Setup for Partial Adherence in Principal Stratification Analysis

For a supposed clinical trial, define the following variables for individual i , and suppose there are a total of n participants in a trial. Note that outcome is presumed to be a binary variable for simplicity, but could be extended to a continuous variable:

$$\begin{aligned}
 Z_i = \text{Treatment assigned} &\in \begin{cases} 0 & \text{Assigned to not receive treatment with probability } 1 - p_Z, \\ 1 & \text{Assigned to receive treatment with probability } p_Z. \end{cases} \\
 X_i = \text{Level of treatment recieved} &\in \begin{cases} N & \text{No treatment received,} \\ P & \text{Partial treatment received,} \\ F & \text{Full treatment received.} \end{cases} \\
 Y_i = \text{Outcome variable} &\in \{0, 1\}.
 \end{aligned}$$

Note that p_Z is held constant across all individuals, as is the standard for randomized clinical trials. It is further supposed that before random treatment assignment occurs, each individual belongs to one of three principal strata defined as follows:

1. Those who will comply fully with either treatment assigned. This group is called the “fully compliant” (**F**) group.
2. Those who will never begin the treatment. This group is called the “never taking” (**N**) group.
3. Those who will, if assigned treatment, only partially comply with the prescribed treatment schedule. If not assigned treatment, they will not receive treatment as directed. This group is called the “partially compliant” (**P**) group.

The principal strata may be referred to as follows:

$$P_i = \text{Principal Stratum} \in \begin{cases} F & \text{Fully Compliant,} \\ P & \text{Partially Compliant,} \\ N & \text{Never Taker.} \end{cases}$$

Temporal patterns will not be observed here. It is assumed that principal strata are determined before randomization and do not shift in response to the course of active treatment or control treatment. The principal stratum of an individual determines the nature of the relationship between Z and X as follows:

$$X_i = \begin{cases} N & \text{if } Z_i = 0 \text{ or } P_i = N, \\ P & \text{if } Z_i = 1 \text{ and } P_i = P, \\ F & \text{if } Z_i = 1 \text{ and } P_i = F. \end{cases}$$

Next, notations will be set up for some values that would be calculated from a full data set of Z_i, X_i, Y_i values collected for individuals $i = 1, \dots, n$. Let n_j = the number of participants

with $Z_i = j$, and let $\bar{Y}_j$ = the proportion of participants with $Z_i = j$ that have outcome $Y_i = 1$. Thus, $\bar{Y}_j$ can be represented as $\frac{1}{n_j} \sum_{i:Z_i=j} Y_i$.

Similarly, let $n_{j,k}$ = the number of participants with $Z_i = j$ and $X_i = k$, and $\bar{Y}_{j,k}$ = the proportion of participants with $Z_i = j$ and $X_i = k$ that have outcome $Y_i = 1$, which can then be represented as $\frac{1}{n_{j,k}} \sum_{i:Z_i=j, X_i=k} Y_i$.

It is thus seen that given a data set, there may be values calculated for n_0, n_1 and $\bar{Y}_0, \bar{Y}_1$. There can also be values calculated for $n_{0,n}, n_{1,n}, n_{1,p}, n_{1,f}$, and $\bar{Y}_{0,n}, \bar{Y}_{1,n}, \bar{Y}_{1,p}, \bar{Y}_{1,f}$. Note that the terms $n_{0,p}$ and $n_{0,f}$ will be 0 while $\bar{Y}_{0,p}$ and $\bar{Y}_{0,f}$ will be undefined if considered, as it is assumed that anyone not assigned treatment can not access treatment, meaning for cases where $Z_i = 0$ it is certain that $X_i = N$ in any observable data.

The goal of the principal stratification approach is to form estimates of average outcomes for treatments within each principal stratum, so notation must be defined for these parameters that are to be estimated.

γ_j = Population proportion of persons with $Z = j$ that will have outcome $Y = 1$.

$\gamma_{k,p}$ = Population proportion of persons with $X = k$ and $P = p$ that will have outcome $Y = 1$.

p_p = True population proportion of persons with $P = p$.

There are two equations which show the relationship of these parameters:

$$\gamma_1 = \gamma_{f,FPF} + \gamma_{p,PPP} + \gamma_{n,NPN}, \quad (1)$$

$$\gamma_0 = \gamma_{n,FPF} + \gamma_{n,PPP} + \gamma_{n,NPN}. \quad (2)$$

A selection of these parameters may be consistently estimated from the data set. First,

$$\hat{\gamma}_1 = \bar{Y}_1,$$

$$\hat{\gamma}_0 = \bar{Y}_0.$$

Next, in observations where $Z = 0$, no information can be inferred about principal stratum, because $X = 0$ always when $Z = 0$. Conversely, observations where $Z = 1$ can have principal stratum fully identified by observing the X value. Thus, consistent estimation of principal strata proportions can be achieved using only the $Z = 1$ data as follows:

$$\begin{aligned}\hat{p}_N &= \frac{n_{1,n}}{n_1}, \\ \hat{p}_P &= \frac{n_{1,p}}{n_1}, \\ \hat{p}_F &= \frac{n_{1,f}}{n_1}.\end{aligned}$$

There are also estimates for cases where the treatment assigned and treatment received correspond to one and only one combination of principal strata and treatment received, in other words when the exact principal stratum of the patient can be identified. These cases inform the following consistent parameter estimates. Specifically, being assigned treatment and receiving full treatment indicates an F patient, being assigned treatment and receiving partial treatment indicates a P patient, and being assigned treatment and receiving no treatment indicates an N patient. As a result:

$$\begin{aligned}\hat{\gamma}_{f,F} &= \bar{Y}_{1,f}, \\ \hat{\gamma}_{p,P} &= \bar{Y}_{1,p}, \\ \hat{\gamma}_{n,N} &= \bar{Y}_{1,n}.\end{aligned}$$

All of the above estimates can be substituted into equations (1) and (2), leaving the following:

$$\hat{\gamma}_1 = \hat{\gamma}_{f,F}\hat{p}_F + \hat{\gamma}_{p,P}\hat{p}_P + \hat{\gamma}_{n,N}\hat{p}_N, \quad (3)$$

$$\hat{\gamma}_0 = \gamma_{n,F}\hat{p}_F + \gamma_{n,P}\hat{p}_P + \gamma_{n,N}\hat{p}_N. \quad (4)$$

Here it can be seen that equation (3) already has all elements estimated, while equation (4)

has two parameters without estimates, meaning consistent estimates for $\gamma_{n,F}$ or $\gamma_{n,P}$ have not been achieved.

2.2 Possible Assumptions and Their Implications

If considering all combinations of treatment received and principal strata, there are 9 parameters defined using the $\gamma_{k,p}$ notation. This is despite only 5 parameters governing the distribution of the observable data: $\gamma_{f,F}, \gamma_{n,F}, \gamma_{p,P}, \gamma_{n,P}, \gamma_{n,N}$. Allowing for the hypothetical existence of the other 4 counterfactual parameters, it is useful to consider them in a table as follows.

Table 1: Mean responses for all treatment and principal stratum combinations, with counterfactual values denoted with grey shading.

		Treatment Received		
		Full	Partial	None
Principal Stratum	F	$\gamma_{f,F}$	$\gamma_{p,F}$	$\gamma_{n,F}$
	P	$\gamma_{f,P}$	$\gamma_{p,P}$	$\gamma_{n,P}$
	N	$\gamma_{f,N}$	$\gamma_{p,N}$	$\gamma_{n,N}$

Two possible assumptions could be made regarding these mean response parameters, and (depending on the situation and validity of the assumptions) they may help in forming more precise estimates of treatment effects. The two possible assumptions are:

Assumption 1: Within a specific treatment, the parameter values are monotone from principal strata N to P to F. The direction of monotonicity is the same for each treatment. This will be otherwise known as the “Monotone Compliance Response” assumption. Referring to Table 1, this corresponds to each column being monotone in the same direction from the top row to the bottom row. Disregarding the counterfactual parameters, this assumption only results in the conclusion that $\min(\gamma_{n,F}, \gamma_{n,N}) \leq \gamma_{n,P} \leq \max(\gamma_{n,F}, \gamma_{n,N})$.

Assumption 2: Within a specific principal stratum, the mean response parameter values are monotone from no treatment to partial treatment to full treatment. The direction of monotonicity is the same in each principal stratum. This will be otherwise known as the “Monotone Dose Response” assumption. Referring to Table 1, this corresponds to each row being monotone in the same direction from the left column to the right column. Disregarding the counterfactual parameters, this assumption only results in the conclusion that $\gamma_{f,F} - \gamma_{n,F}$ shares its sign with $\gamma_{p,P} - \gamma_{n,P}$.

2.2.1 Bounds Resulting From the Monotone-within-Treatment Assumption

Taking Assumption 1 (monotone-within-treatment) to be true results in the following:

$$\min(\gamma_{n,F}, \gamma_{n,N}) \leq \gamma_{n,P} \leq \max(\gamma_{n,F}, \gamma_{n,N}).$$

To see in what way this bounds estimation of $\gamma_{n,F}$ and $\gamma_{n,P}$, Equation (2) can be used to rearrange and express $\gamma_{n,F}$ in terms of other parameters as follows:

$$\gamma_{n,F} = \frac{\gamma_0 - \gamma_{n,PPP} - \gamma_{n,NPN}}{p_F}.$$

The following is then arrived at, from which the situation branches into two cases:

$$\min\left(\frac{\gamma_0 - \gamma_{n,PPP} - \gamma_{n,NPN}}{p_F}, \gamma_{n,N}\right) \leq \gamma_{n,P} \leq \max\left(\frac{\gamma_0 - \gamma_{n,PPP} - \gamma_{n,NPN}}{p_F}, \gamma_{n,N}\right). \quad (5)$$

It is uncertain which is larger in the minimum/maximum functions, so it is tentatively assumed that one case is true and the conclusion is found, then assumed the other case and the same process is completed. Then, the two conclusions can be combined into the most general answer.

Case 1, $\gamma_{n,N}$ is the minimum in Equation (5)

The two sides of the resulting inequality are separately rearranged, so as to have $\gamma_{n,P}$ in the center, because it is the only non-estimated quantity present.

$$\begin{aligned}
\gamma_{n,N} &\leq \gamma_{n,P} & \gamma_{n,P} &\leq \frac{\gamma_0 - \gamma_{n,P}p_P - \gamma_{n,N}p_N}{p_F} \\
\gamma_{n,N} &\leq \gamma_{n,P} & p_F\gamma_{n,P} &\leq \gamma_0 - \gamma_{n,P}p_P - \gamma_{n,N}p_N \\
\gamma_{n,N} &\leq \gamma_{n,P} & (p_F + p_P)\gamma_{n,P} &\leq \gamma_0 - \gamma_{n,N}p_N \\
\gamma_{n,N} &\leq \gamma_{n,P} & \gamma_{n,P} &\leq \frac{\gamma_0 - \gamma_{n,N}p_N}{p_F + p_P}
\end{aligned}$$

An inequality that provides information for the estimation of $\gamma_{n,P}$ is now directly observable, and $\gamma_{n,F}$ may be substituted in terms of the other terms in Equation (2) to also provide information for the estimation of $\gamma_{n,F}$. Thus the following two intervals are seen for the two unknown parameters $\gamma_{n,F}$ and $\gamma_{n,P}$, after some algebra

$$\begin{aligned}
\gamma_{n,N} &\leq \gamma_{n,P} \leq \frac{1}{p_F + p_P}\gamma_0 - \frac{p_N}{p_F + p_P}\gamma_{n,N}, \\
\frac{1}{p_F + p_P}\gamma_0 - \frac{p_N}{p_F + p_P}\gamma_{n,N} &\leq \gamma_{n,F} \leq \frac{1}{p_F}\gamma_0 - \frac{p_P + p_N}{p_F}\gamma_{n,N}.
\end{aligned}$$

Instead of having bounds for $\gamma_{n,F}$ and $\gamma_{n,P}$, it might be desired to have bounds for the actual treatment effects within the principal strata: $(\gamma_{f,F} - \gamma_{n,F})$ and $(\gamma_{p,P} - \gamma_{n,P})$. These would be expressed as follows:

$$\begin{aligned}
\gamma_{p,P} - \frac{1}{p_F + p_P}\gamma_0 + \frac{p_N}{p_F + p_P}\gamma_{n,N} &\leq \gamma_{p,P} - \gamma_{n,P} \leq \gamma_{p,P} - \gamma_{n,N}, \\
\gamma_{f,F} - \frac{1}{p_F}\gamma_0 + \frac{p_P + p_N}{p_F}\gamma_{n,N} &\leq \gamma_{f,F} - \gamma_{n,F} \leq \gamma_{f,F} - \frac{1}{p_F + p_P}\gamma_0 + \frac{p_N}{p_F + p_P}\gamma_{n,N}.
\end{aligned}$$

Case 2, $\gamma_{n,N}$ is the maximum in Equation (5)

Again, steps similar to the previous case will be used to separately rearrange the two sides of the resulting inequality, so as to have $\gamma_{n,P}$ in the center, because it is the only non-estimable

quantity present. After performing similar arithmetic as in the previous case, the following intervals are arrived at for $\gamma_{n,F}$ and $\gamma_{n,P}$:

$$\begin{aligned}\frac{1}{p_F + p_P}\gamma_0 - \frac{p_N}{p_F + p_P}\gamma_{n,N} &\leq \gamma_{n,P} \leq \gamma_{n,N}, \\ \frac{1}{p_F}\gamma_0 - \frac{p_P + p_N}{p_F}\gamma_{n,N} &\leq \gamma_{n,F} \leq \frac{1}{p_F + p_P}\gamma_0 - \frac{p_N}{p_F + p_P}\gamma_{n,N}.\end{aligned}$$

And thus the intervals from Case 1 have been flipped in Case 2 for the estimated $\gamma_{n,F}$ and $\gamma_{n,P}$ values. Similarly, extending to interval estimates of $(\gamma_{f,F} - \gamma_{n,F})$ and $(\gamma_{p,P} - \gamma_{n,P})$, flipped intervals are seen once more:

$$\begin{aligned}\gamma_{p,P} - \gamma_{n,N} &\leq \gamma_{p,P} - \gamma_{n,P} \leq \gamma_{p,P} - \frac{1}{p_F + p_P}\gamma_0 + \frac{p_N}{p_F + p_P}\gamma_{n,N}, \\ \gamma_{f,F} - \frac{1}{p_F + p_P}\gamma_0 + \frac{p_N}{p_F + p_P}\gamma_{n,N} &\leq \gamma_{f,F} - \gamma_{n,F} \leq \gamma_{f,F} - \frac{1}{p_F}\gamma_0 + \frac{p_P + p_N}{p_F}\gamma_{n,N}.\end{aligned}$$

Combining the two cases

Without ever trying to infer which terms were the smallest or largest in the original minimum/maximum functions, the desired intervals may simply be expressed as being between the minimum of the two bounding terms and the maximum of the two bounding terms, since the two cases simply have flipped intervals.

In other words, if the following are defined:

$$\begin{aligned}A &= \gamma_{p,P} - \gamma_{n,N}, \\ B &= \gamma_{p,P} - \frac{1}{p_F + p_P}\gamma_0 + \frac{p_N}{p_F + p_P}\gamma_{n,N}, \\ C &= \gamma_{f,F} - \frac{1}{p_F + p_P}\gamma_0 + \frac{p_N}{p_F + p_P}\gamma_{n,N}, \\ D &= \gamma_{f,F} - \frac{1}{p_F}\gamma_0 + \frac{p_P + p_N}{p_F}\gamma_{n,N}.\end{aligned}$$

Then it follows that:

$$\begin{aligned}\min(A, B) &\leq \gamma_{p,P} - \gamma_{n,P} \leq \max(A, B), \\ \min(C, D) &\leq \gamma_{f,F} - \gamma_{n,F} \leq \max(C, D).\end{aligned}$$

It is clear that A, B, C, D may be estimated consistently, and therefore there are consistently estimable bounds by using the estimates $\hat{\gamma}_{p,P}, \hat{\gamma}_{n,N}, \hat{\gamma}_{f,F}, \hat{p}_P, \hat{p}_N, \hat{p}_F, \hat{\gamma}_0$ to obtain $\hat{A}, \hat{B}, \hat{C}, \hat{D}$.

2.3 A Monte Carlo Method of Generating Posterior Samples of Effect Estimates

It has been seen that introducing a third treatment level to distinguish partial adherers to study protocols results in an ability to form consistent parameter estimates using methods similar to a simpler principal stratification approach, but only for some parameters. For other parameters, it has been seen that only consistent bound estimates may be computed. As a result, one next step in developing this expansion to the principal stratification approach is to consider the possibility of determining posterior distributions of effects of interest.

The present model for the data includes unknown parameters for how the population is distributed across principal strata (p_B, p_P, p_N) , as well as unknown parameters for probabilities of the outcome of interest for each treatment type within each principal strata (9 different $\gamma_{Q,W}$ values, for treatment levels Q from f, p, n and principal strata W from B, P, N). Out of these parameters, only some values and linear combinations of values may be consistently estimated with increasing accuracy as the study sample size increases, while others will not become informed by gathering data in the given study framework. For this reason, the model is called partially identified. By following the steps outlined in Gustafson [21] for partially identified models as described below, a Monte Carlo simulation method can be designed that is computationally less expensive than other methods such as Markov Chain Monte Carlo:

1. Identify a ‘transparent reparameterization’: a one-to-one invertible mapping h from

the original parameter set θ to a new parameter set defined such that it can be strictly separated into fully identifiable (ϕ) and non-identifiable (λ) parts.

2. Transform a specified prior distribution $\pi_\theta(\theta)$ such that it is expressed in terms of the new parameterization.
3. Identify a convenience prior distribution in the new parameter space $\pi^*(\phi, \lambda)$. We will have that $\pi^*(\phi, \lambda)$ is independently uniform, and it follows that:

$$\pi_{\phi, \lambda|y}^*(\phi, \lambda|y) = \pi^*(\phi|y)\pi^*(\lambda|\phi)$$

$$\pi_{\phi, \lambda|y}^*(\phi, \lambda|y) \propto \pi^*(\phi|y)$$

$$\pi_{\phi, \lambda|y}^*(\phi, \lambda|y) \propto L(y|\phi)\pi^*(\phi)$$

$$\pi_{\phi, \lambda|y}^*(\phi, \lambda|y) \propto L(y|\phi)$$

4. Determine the proportional estimator of the convenience posterior $L(y|\phi)$ by separating it into likelihood components. The y vector can be separated into four different groups of independent and assumed identically distributed observations. Thus the target likelihood can be calculated as the product of each of the four groups' outcome likelihoods as well as the likelihood of the groups' distribution in the population.
5. An importance sampling scheme [34] is used to sample from the convenience posterior distribution. Weights are assigned to each sample in order to adjust from convenience posterior samples to samples representing the target posterior distribution.

Step 1: Transparent Reparameterization

Let the original parameter set be defined as the vector θ :

$$\theta = (\gamma_{n,P} \quad \gamma_{f,F} \quad \gamma_{p,P} \quad \gamma_{n,N} \quad p_F \quad p_P \quad p_N \quad \gamma_{n,F} \quad \gamma_{p,F} \quad \gamma_{f,P} \quad \gamma_{f,N} \quad \gamma_{p,N})^T.$$

Referring back to the derivations directly preceding Equations (3) and (4), it has been shown that consistent estimates are available for $\gamma_{f,F}, \gamma_{p,P}, \gamma_{n,N}, p_F, p_P$, and p_N . Furthermore, a consistent estimate is available for $\gamma_{n,F}p_F + \gamma_{n,P}p_P + \gamma_{n,N}p_N$, as it is equivalent to γ_0 , the average outcomes seen in those assigned control intervention, as per Equation (2). Define the following parameter for simplicity in further derivations:

$$\chi = \gamma_{n,F}p_F + \gamma_{n,P}p_P + \gamma_{n,N}p_N.$$

With this information, it is possible to form a simple mapping from θ to a set of parameters that clearly distinguish parameters fully identifiable by data from parameters that are not at all identifiable by data. Calling this mapping h , the new set of identifiable parameters ϕ , and the new set of non-identifiable parameters λ , the mapping is written as follows:

$$h(\theta) = h \begin{pmatrix} \gamma_{n,P} \\ \gamma_{f,F} \\ \gamma_{p,P} \\ \gamma_{n,N} \\ p_F \\ p_P \\ p_N \\ \gamma_{n,F} \\ \gamma_{p,F} \\ \gamma_{f,P} \\ \gamma_{f,N} \\ \gamma_{p,N} \end{pmatrix} = \begin{pmatrix} \chi \\ \gamma_{f,F} \\ \gamma_{p,P} \\ \gamma_{n,N} \\ p_F \\ p_P \\ p_N \\ \gamma_{n,F} \\ \gamma_{p,F} \\ \gamma_{f,P} \\ \gamma_{f,N} \\ \gamma_{p,N} \end{pmatrix} = \left(\frac{\phi}{\lambda} \right) = (\phi, \lambda)^T.$$

In the above mapping, the resulting vector has a horizontal line which shows the distinction between the parameters fully informed by data (ϕ) and the parameters not informed by data (λ). It can be observed that this transparent reparameterization is simple in that

the mapping h maps most parameters to themselves, and only one parameter has been mapping to a linear combination of itself and some other parameters, as $\gamma_{n,P}$ maps to $\chi = \gamma_{n,FPB} + \gamma_{n,PPP} + \gamma_{n,NPN}$.

Step 2: Target Prior

We specify a prior on the original parameter space θ as independently Uniform(0,1) for each $\gamma_{Q,W}$, and Dirichlet(1,1,1) for (p_B, p_P, p_N) . Then, in order to express the prior distribution in terms of the new parameters (ϕ, λ) , the following identity may be used:

$$\pi_{\phi,\lambda}(\phi, \lambda) = \pi_{\theta}(h^{-1}(\phi, \lambda)) \left| \frac{\partial(h^{-1})}{\partial(\phi, \lambda)} \right|.$$

Due to the fact that $\pi_{\theta}(\theta)$ is uninformative except for defining the bounds of the parameter space, the term $\pi_{\theta}(h^{-1}(\phi, \lambda))$ will simply be an indicator function that maps to 0 if the inputs (ϕ, λ) lie outside the image of h . Let $\mathbb{H}$ represent the image of h . The determinant of the inverse mapping's derivative is simpler than it may appear, because of the fact that the mapping h and thus h^{-1} maps most vector elements to themselves. In the inverse derivative matrix, it arises that all but the first row and first column are as in an identity matrix, and thus the determinant of the matrix will simply be the element in the (1,1) position of the matrix, which is the derivative of the first element of h^{-1} (the mapping from χ back to $\frac{\chi - \gamma_{n,FPB} - \gamma_{n,NPN}}{p_P}$) in terms of χ , which is simply p_P^{-1} . Thus, we arrive at the specified prior in terms of the new parameterization as follows:

$$\pi_{\phi,\lambda}(\phi, \lambda) = \frac{I_{\mathbb{H}}(\phi, \lambda)}{p_P}.$$

Step 3: Convenience Prior

For more convenient calculations, it is necessary to identify a convenience prior in the new parameter space. Similar to the true prior in the original parameter space, define the con-

venience prior $\pi^*(\phi, \lambda)$ to be independently $\text{Unif}(0,1)$ for χ and the $\gamma_{Q,W}$ parameters, and $\text{Dirichlet}(1,1,1)$ for p_B, p_P, p_N . Then, due to the independence of the parameters, it is true that $\pi^*(\lambda|\phi)$ is equal to $\pi^*(\lambda)$ and thus uniformly distributed, and it is true that $\pi^*(\phi)$ is also uniform.

Step 4: Convenience Posterior

Due to the way the convenience prior was defined, the following simplifying steps can be used when forming an expression for the convenience posterior distribution:

$$\begin{aligned}\pi^*(\phi, \lambda|y) &= \pi^*(\phi|y)\pi^*(\lambda|\phi) \\ &\propto \pi^*(\phi|y) \\ &\propto L(y|\phi)\pi^*(\phi) \\ &\propto L(y|\phi)\end{aligned}$$

So it is seen that when the likelihood of the data given ϕ is calculated, it shall be proportional to the convenience posterior distribution. At this point the likelihood function can be broken into 5 components. There are four separable ‘bins’ that independent observations can belong to: one for those assigned no treatment, then three bins for those assigned treatment for the three possible treatments received. Each of these bins can independently have its likelihood calculated using relevant terms. The fifth component of the likelihood function is the likelihood that the observations were assigned to their respective bins. Subscripts are used here to indicate specific subsets of observations (e.g. $y_{z=1,x=0}$ is the set of outcomes only corresponding to observations with $z = 1$ and $x = 0$) or to indicate counts of observations (e.g. $n_{z=1,x=0}$ is the number of observations with $z = 1, x = 0$). The likelihood function’s separation into components can be expressed as follows:

$$L(y|\phi) = L(y_{z=0}|\phi)L(y_{z=1,x=0}|\phi)L(y_{z=1,x=1}|\phi)L(y_{z=1,x=2}|\phi)L(n_{z=0}, n_{z=1,x=0}, n_{z=1,x=2}, n_{z=1,x=2}|\phi).$$

And here, each part of the likelihood function now only depends on certain parts of ϕ :

$$L(y|\phi) = L(y_{z=0}|\chi)L(y_{z=1,x=0}|\gamma_{n,N})L(y_{z=1,x=1}|\gamma_{p,P})L(y_{z=1,x=2}|\gamma_{f,F}) \times \\ L(n_{z=0}, n_{z=1,x=0}, n_{z=1,x=2}, n_{z=1,x=2}|p_F, p_P, p_N).$$

And now, the first four likelihoods can be expressed using simple binomial distribution functions, while the fifth likelihood does not vary with $n_{z=0}$, since no parameters inform this value, but the remaining three n values may be expressed in a multinomial distribution. Disregarding terms not involving the parameters, the resulting form is:

$$L(y|\phi) \propto \chi^{n_{z=0,y=1}}(1 - \chi)^{n_{z=0,y=0}}\gamma_{n,N}^{n_{z=1,x=0,y=1}}(1 - \gamma_{n,N})^{n_{z=1,x=0,y=0}}\gamma_{p,P}^{n_{z=1,x=1,y=1}}(1 - \gamma_{p,P})^{n_{z=1,x=1,y=0}} \\ \gamma_{f,F}^{n_{z=1,x=2,y=1}}(1 - \gamma_{f,F})^{n_{z=1,x=2,y=0}}p_N^{n_{z=1,x=0}}p_P^{n_{z=1,x=1}}p_F^{n_{z=1,x=2}}.$$

From this, the convenience posterior $\pi^*(\phi, \lambda|y)$ can be identified as follows, where all parameters have posterior distributions independent of each other:

$$\begin{aligned} \chi &\sim B(n_{z=0,y=1} + 1, n_{z=0,y=0} + 1). \\ \gamma_{f,F} &\sim B(n_{z=1,x=2,y=1} + 1, n_{z=1,x=2,y=0} + 1). \\ \gamma_{p,P} &\sim B(n_{z=1,x=1,y=1} + 1, n_{z=1,x=1,y=0} + 1). \\ \gamma_{n,N} &\sim B(n_{z=1,x=0,y=1} + 1, n_{z=1,x=0,y=0} + 1). \\ (p_F, p_P, p_N) &\sim \text{Dirichlet}(n_{z=1,x=2} + 1, n_{z=1,x=1} + 1, n_{z=1,x=0} + 1). \\ \gamma_{n,F} &\sim \text{Unif}(0, 1). \\ \gamma_{p,F} &\sim \text{Unif}(0, 1). \\ \gamma_{f,P} &\sim \text{Unif}(0, 1). \\ \gamma_{f,N} &\sim \text{Unif}(0, 1). \\ \gamma_{p,N} &\sim \text{Unif}(0, 1). \end{aligned}$$

Step 5: Target Posterior Sampling Scheme

In order to adjust any samples taken from the convenience posterior to reflect the true posterior distribution, weights are assigned to each sample taken from the convenience posterior that are defined such that the following hold. Letting m represent the size of the Monte Carlo sample, we take:

$$w_i \propto \frac{\pi(\phi_i, \lambda_i)}{\pi^*(\phi_i, \lambda_i)},$$

With

$$\sum_{i=1}^m w_i = 1.$$

As $\pi(\phi_i, \lambda_i)$ and $\pi^*(\phi_i, \lambda_i)$ have already been determined, it is seen that:

$$w_i \propto \frac{I_{\mathbb{H}}(\phi, \lambda)}{p_P}.$$

In the following section, it will be more clearly shown how to use these results (particularly the posterior distributions from Step 4 and the weights from Step 5) in order to build a posterior sampling scheme.

3 Examples and Simulation Results

In the following sections, the “Monte Carlo method” refers to the posterior sampling method previously formulated in Section 2.3. One recurring choice that will be discussed for all simulations is the choice of whether any of the monotonicity types discussed in Section 2.2 are satisfied during data simulation, and furthermore the choice of whether any of the same monotonicity types are assumed to be true when performing inference on simulated data sets. In the following sections, the phrase “monotonicity type X” refers to the case where

“assumption X” is true during data simulation, while saying “assumption X” only refers to using the assumption during inference and not during data simulation.

The main idea in the setup of the forthcoming simulations is that different monotonicity types can be crossed with different assumptions. For example, the case of data being sampled from a distribution satisfying monotonicity type 1 and being analyzed using assumption 2 could potentially be very different from the case of data being sampled from a distribution satisfying monotonicity types 1 and 2 and being analyzed using no assumptions. As a result, forthcoming simulations try to observe the full scope of combinations of monotonicity types used when generating data and of assumptions made when analyzing data, although only a selection of these combinations are shown in full detail.

The steps taken to generate data using a specific monotonicity type is as follows:

1. Decide on a parameter set θ (as defined in Section 2.3) that satisfies the given monotonicity type(s). Also, decide on n , the number of people in the simulated study, and p , the proportion of patients who would be randomly assigned treatment (for all simulations shown in this thesis, $p = 0.5$).
2. Randomly generate the vector of Z_i values as n samples from a Bernoulli(p) distribution.
3. Randomly generate principal strata P_i by sampling n values from the three possible principal strata with replacement, with sampling weights p_F, p_P, p_N as determined in θ .
4. Determine the vector of X_i values using the P_i and Z_i values (this may be coded using nested ‘if’ statements, such that for example if $Z_i = 1$ and $P_i = F$, then $X_i = F$).
5. Define a new vector of probabilities K_i that the i^{th} patient will have outcome $Y_i = 1$

based on θ , P_i and X_i (this may again be coded using nested ‘if’ statements, such that for example if $P_i = F$ and $X_i = F$, then $K_i = \gamma_{f,F}$).

6. Generate the n values of Y_i from Bernoulli(K_i) distributions.
7. The generated data to be analyzed are the vectors of Z_i , X_i , and Y_i values.

After simulating data, posterior estimates $\hat{\theta}$ are generated using the Monte Carlo method, which is done using the following steps, drawing on the results from Section 2.3:

1. Select which assumption(s) will be made for the analysis, if any. Select the desired posterior sample size, N .
2. Define some G such that $G \geq N$ (some samples will be assigned weights of 0 as a result violating requirements of posterior samples, reducing the actual sample posterior size, so more should be generated to accommodate this). We generally used $G = 5N$.
3. Using the data set of Z_i , X_i , and Y_i values, sample G values from the following distributions as (unweighted) posterior samples for each parameter:

$$\chi \sim B(n_{z=0,y=1} + 1, n_{z=0,y=0} + 1).$$

$$\gamma_{f,F} \sim B(n_{z=1,x=2,y=1} + 1, n_{z=1,x=2,y=0} + 1).$$

$$\gamma_{p,P} \sim B(n_{z=1,x=1,y=1} + 1, n_{z=1,x=1,y=0} + 1).$$

$$\gamma_{n,N} \sim B(n_{z=1,x=0,y=1} + 1, n_{z=1,x=0,y=0} + 1).$$

$$(p_F, p_P, p_N) \sim \text{Dirichlet}(n_{z=1,x=2} + 1, n_{z=1,x=1} + 1, n_{z=1,x=0} + 1).$$

$$\gamma_{n,F} \sim \text{Unif}(0, 1).$$

$$\gamma_{p,F} \sim \text{Unif}(0, 1).$$

$$\gamma_{f,P} \sim \text{Unif}(0, 1).$$

$$\gamma_{f,N} \sim \text{Unif}(0, 1).$$

$$\gamma_{p,N} \sim \text{Unif}(0, 1).$$

4. Calculate the posterior sample $\gamma_{n,P}$ values as $(\chi - \gamma_{n,N}p_N - \gamma_{n,F}p_F)p_P^{-1}$.
5. Identify and remove any posterior observations that violate assumptions, and any posterior observations that have $\gamma_{n,P}$ estimates less than 0 or greater than 1. If the total number of remaining samples is less than N , repeat the previous two steps. This will continue until at least N posterior samples are achieved.
6. Calculate weights such that each w_i is proportional to the inverse of the i^{th} posterior observation of p_P , and the sum of all weights is 1.
7. The posterior samples are thus the remaining samples for each parameter, with their associated weights w_i .

Forthcoming figures will only show posterior distributions of three parameters of particular interest: the effect of receiving partial treatment for partial adherers, the effect of receiving full treatment for full adherers, and the average effect of receiving full treatment within the whole population. These three are expressed algebraically in their respective order as:

$$\psi_1 = \gamma_{p,P} - \gamma_{n,P},$$

$$\psi_2 = \gamma_{f,F} - \gamma_{n,F},$$

$$\psi_3 = (\gamma_{f,FPB} + \gamma_{f,PPP} + \gamma_{f,NPN}) - (\gamma_{n,FPB} + \gamma_{n,PPP} + \gamma_{n,NPN}).$$

3.1 Limiting Distributions of Parameter Estimates

As described in section 2.2.1, when Assumption 1 is made, data may be used to estimate $\hat{A}$, $\hat{B}$, $\hat{C}$, and $\hat{D}$, and these values may be used to form estimated bounds on two quantities of interest: the average effect of partial treatment in partially compliant people $(\gamma_{p,P} - \gamma_{n,P})$ and the average effect of full treatment in fully compliant people $(\gamma_{f,F} - \gamma_{n,F})$. As n tends towards infinity these estimated bounds will converge to true bounds. Just as the bounds converge to true bounds as study sample size approaches infinity, so too will the posterior distributions

of target parameters converge to limiting posterior distributions. It is of interest to visualize using some examples of how estimated posterior distributions and estimated bounds compare to limiting bounds and the limiting posterior distribution. Two examples will be visualized, where one has data generated from a distribution satisfying monotonicity types 1 and 2 and thus no invalid assumptions, and one has no monotonicity types satisfied, and thus neither of the assumptions are valid. Both have a study size of 1000 for the estimated bounds and distribution, and a study size approaching infinity for the limiting bounds and distribution.

3.1.1 Simulation 1: Both Monotonicity Types

The parameter set used for the case with both monotonicity types is:

$$\theta_{1,2} = \begin{pmatrix} \gamma_{n,P} \\ \gamma_{f,F} \\ \gamma_{p,P} \\ \gamma_{n,N} \\ p_F \\ p_P \\ p_N \\ \gamma_{n,F} \\ \gamma_{p,F} \\ \gamma_{f,P} \\ \gamma_{f,N} \\ \gamma_{p,N} \end{pmatrix} = \begin{pmatrix} 0.4 \\ 0.55 \\ 0.45 \\ 0.35 \\ 0.7 \\ 0.2 \\ 0.1 \\ 0.45 \\ 0.5 \\ 0.5 \\ 0.45 \\ 0.4 \end{pmatrix}.$$

Following in Figures 1, 2, and 3 are plots showing posterior sample values for the parameters of interest ψ_1, ψ_2 , and ψ_3 respectively. For each possible assumption, the posterior sample distribution for a study of size 1000 is compared to the posterior sample distribution for a study of a size approaching infinity (with bounds additionally shown when Assumption 1 is

made).

Figure 1: Histogram of the posterior sample of the partial treatment effect on P patients, with monotonicity types 1 and 2 satisfied. Posteriors formed using violated assumptions have a grey fill, whereas samples with all assumptions valid have a white fill. A solid vertical line indicates the true value, dashed indicates estimated bounds, dotted indicates limiting bounds. The limiting posterior density curve is plotted over the histogram.

Figure 2: Histogram of the posterior sample of the full treatment effect on F patients, with monotonicity types 1 and 2 satisfied. Posteriors formed using violated assumptions have a grey fill, whereas samples with all assumptions valid have a white fill. A solid vertical line indicates the true value, dashed indicates estimated bounds, dotted indicates limiting bounds. The limiting posterior density curve is plotted over the histogram.

Figure 3: Histogram of the posterior sample of the full treatment effect on all patients with monotonicity types 1 and 2 satisfied. Posteriors formed using violated assumptions have a grey fill, whereas samples with all assumptions valid have a white fill. A solid vertical line indicates the true value. The limiting posterior density curve is plotted over the histogram.

In Figures 1, 2, and 3 it can be seen that invoking Assumption 2 does not affect inference as much as Assumption 1 in this example, where Assumption 1 appears to reduce variation in the posterior distribution of the partial treatment effect on PC patients and the full treatment effect on BC patients, but less so in the full treatment effect on the population. Furthermore, it is visible that when assumptions were made that allowed for bounds to be calculated, there were cases where estimated bounds completely missed the true value, but as sample size increased these estimated bounds converged to accurate ones. It is also worth noting that estimated bounds seem to agree with posterior, as they existed somewhat near the modes of estimates in these examples.

3.1.2 Simulation 2: Neither Monotonicity Type

A parameter set now used for the case where no monotonicity types are actually satisfied is:

$$\theta_0 = \begin{pmatrix} \gamma_{n,P} \\ \gamma_{f,F} \\ \gamma_{p,P} \\ \gamma_{n,N} \\ p_F \\ p_P \\ p_N \\ \gamma_{n,F} \\ \gamma_{p,F} \\ \gamma_{f,P} \\ \gamma_{f,N} \\ \gamma_{p,N} \end{pmatrix} = \begin{pmatrix} 0.5 \\ 0.55 \\ 0.45 \\ 0.55 \\ 0.7 \\ 0.2 \\ 0.1 \\ 0.55 \\ 0.5 \\ 0.5 \\ 0.55 \\ 0.5 \end{pmatrix}.$$

In the same format as in Simulation 1, following are Figures 4, 5, and 6 showing posterior sample values for the parameters of interest ψ_1, ψ_2 , and ψ_3 respectively. Again, for each possible assumption, the posterior sample distribution for a study of size 1000 is compared to the posterior sample distribution for a study with size approaching infinity (with bounds additionally shown when Assumption 1 is made).

Figure 4: Histogram of the posterior sample of partial treatment effect on P patients, with no monotonicity types satisfied. Posteriors formed using violated assumptions have a grey fill, whereas samples with all assumptions valid have a white fill. A solid vertical line indicates the true value, dashed indicates estimated bounds, dotted indicates limiting bounds. The limiting posterior density curve is plotted over the histogram when Assumption 1 is not made, and omitted when Assumption 1 is made due to the height of the spiked density within the bounds of the dotted lines.

Figure 5: Histogram of the posterior sample of the full treatment effects on F patients, with no monotonicity types satisfied. Posteriors formed using violated assumptions have a grey fill, whereas samples with all assumptions valid have a white fill. A solid vertical line indicates the true value, dashed indicates estimated bounds, dotted indicates limiting bounds. The limiting posterior density curve is plotted over the histogram when Assumption 1 is not made, and omitted when Assumption 1 is made due to the height of the spiked density within the bounds of the dotted lines.

Figure 6: Histogram of the posterior sample of the full treatment effects on all patients, with no monotonicity types satisfied. Posteriors formed using violated assumptions have a grey fill, whereas samples with all assumptions valid have a white fill. A solid vertical line indicates the true value. The limiting posterior density curve is plotted over the histogram.

In Figures 4 and 5, it can be seen that invoking Assumption 1 incorrectly has caused posterior densities to converge to an incorrect interval when estimating the treatment effect among F or P patients, while invoking Assumption 2 incorrectly tended to simply reduce posterior density in the area of the true parameter value compared to surrounding values. In the case of Assumption 2 being made, it seems possible that the posterior mean may still be close to the true value at least in this example, even though the posterior density surrounding the true value has decreased. As in the case where the Assumptions were correctly made in Figures 1 and 2, initial bound estimates can be expected to sometimes miss the true value, but in this case we observe that the limiting bounds do not include the true value.

Invoking incorrect assumptions had different effects on population-level estimates as seen in Figure 6, where incorrectly taking Assumption 2 had a minimal visible effect on posterior probabilities, but taking Assumption 1 with or without Assumption 2 led to the posterior

densities being warped away from the true value, although to a lesser extent than in the cases observed in 4 and 5 with the F and P treatment effects.

Not shown here are the resulting figures corresponding to a data set generated where monotonicity type 1 but not 2 is satisfied, and the resulting figures corresponding to a data set generated where monotonicity type 2 but not 1 is satisfied. The most consistent observations across all completed simulations are that Assumption 1 appears to have a greater impact on inference compared to Assumption 2 – in cases where Assumption 1 is valid, it improves estimate intervals by a notable amount, but in cases where Assumption 1 is invalid, estimate intervals can become entirely incorrect and narrowly focussed on an irrelevant region. In comparison, the effect of Assumption 2 appears more benign, in that when it is incorrectly made, the impact on inference tends to be only slightly away from the true value, while when the assumption is correctly made, it tends to have a small beneficial effect on inference.

3.2 Validation Trial of the Monte Carlo Method through Comparison with JAGS

In order to validate the Monte Carlo simulation method that has been derived, posterior distribution samples can be formed from a simulated data set, comparing the results from using the Monte Carlo simulation method and Markov Chain Monte Carlo sampling implemented using the ‘R2jags’ package in R. Posterior samples are generated for comparison in sixteen cases satisfying all parameter set and inference types, and in all cases posterior samples indicated the same posterior distribution of parameters when comparing the Monte Carlo approach and Markov Chain Monte Carlo sampling. Figures 7, 8, and 9 display posterior sample distributions of all samples taken for the cases where monotonicity type 1 is satisfied. Each figure compares the Monte Carlo method to the MCMC method for all four assumptions, for one estimand of interest.

For this case with monotonicity type 1 satisfied, a study sample size $n = 10000$ is used, with a probability of being assigned treatment of 0.5, and data are simulated using the following parameter set:

$$\theta_2 = \begin{pmatrix} \gamma_{n,P} \\ \gamma_{f,F} \\ \gamma_{p,P} \\ \gamma_{n,N} \\ p_F \\ p_P \\ p_N \\ \gamma_{n,F} \\ \gamma_{p,F} \\ \gamma_{f,P} \\ \gamma_{f,N} \\ \gamma_{p,N} \end{pmatrix} = \begin{pmatrix} 0.5 \\ 0.55 \\ 0.45 \\ 0.45 \\ 0.7 \\ 0.2 \\ 0.1 \\ 0.55 \\ 0.5 \\ 0.5 \\ 0.45 \\ 0.4 \end{pmatrix}.$$

Figure 7: Posterior sample of partial treatment effects on P patients with monotonicity type 2 satisfied. Samples that were taken using violated assumptions have a black fill, whereas samples with all assumptions valid have a white fill.

Figure 8: Posterior sample of full treatment effects on F patients With monotonicity type 2 satisfied. Samples that were taken using violated assumptions have a black fill, whereas samples with all assumptions valid have a white fill.

Figure 9: Posterior sample of full treatment effects on all patients with monotonicity type 2 satisfied. Samples taken using violated assumptions have a black fill, whereas samples with all assumptions valid have a white fill.

The plots for simulated data with no monotonicity, monotonicity type 2, and monotonicity types 1 and 2 (not shown here) have similar results to above. All simulated posterior estimates from the Monte Carlo method and Markov Chain Monte Carlo methods have distributions that are visually indistinguishable as in Figures 7, 8, and 9. The observations are thus consistent with the samples arising from the same distribution.

Since the Monte Carlo method uses an approach comparable to importance sampling, by assigning weights to each observation, it is useful to quantify the effective sample size of the posterior distribution. Kish’s effective sample size [33] is calculated as $(\sum_{i=1}^n w_i^2)^{-1}$, and for all simulations completed with a posterior sample size of over 1000, none of the calculated Kish’s effective sample sizes were more than 1 less than the raw sample sizes. This supports the idea that the weighting used did not severely limit the reliability of the posterior sample.

3.3 Simulation Set Comparing to the Case Discussed by Shrier

The simulations presented by Shrier et. al.[47] were for cases where the outcomes measured were continuous values between 0 and 100, and patients who were assigned to the control group could access treatment (i.e. there could be a further principal stratum of ‘always takers’ in the population). The simulations showed a bias away from the null hypothesis whenever partial treatment was assumed to be full treatment, or when partial treatment was assumed to be no treatment. After having developed the new methods in Sections 2 and 3, it is of interest how these new estimators may compare to the biased approaches observed in Shrier’s simulations.

Thus, we set up a simulation scheme that is as similar to Shrier’s as possible, and is able to compare our methods to the previous biased methods [47]. Instead of having outcomes sampled from normal distributions with a parameter set of means between 0 and 100, we divide these means by 100 and have the parameters instead convey the probability of a binary outcome. Furthermore, all patients in the original simulation scheme that were considered ‘always takers’ are now baseline compliers, so instead of simulating data on 400 full compliers, 200 partial compliers, 200 always takers, and 200 never takers, there are 600 full compliers, 200 partial compliers, and 200 always takers. The simulation steps are as follows:

1. The vector of principal strata P has 600 entries F , 200 entries P , and 200 entries N .
2. The vector of assigned treatment Z satisfies that exactly 300 people with principal stratum F are assigned treatment, 100 people with principal stratum P , and 100 people with principal stratum N .
3. Determine the vector X using the vectors P and Z , which may be coded using nested ‘if’ statements.

4. Use the following parameter vector θ as adapted from Shrier's parameter set to define the vector of probabilities K_i that each patient will have outcome $Y_i = 1$:

$$\theta = \begin{pmatrix} \gamma_{n,P} \\ \gamma_{f,F} \\ \gamma_{p,P} \\ \gamma_{n,N} \\ p_F \\ p_P \\ p_N \\ \gamma_{n,F} \\ \gamma_{p,F} \\ \gamma_{f,P} \\ \gamma_{f,N} \\ \gamma_{p,N} \end{pmatrix} = \begin{pmatrix} 0.7 \\ 0.4 \\ 0.55 \\ 0.9 \\ 0.6 \\ 0.2 \\ 0.2 \\ 0.7 \\ 0.5 \\ 0.45 \\ 0.6 \\ 0.7 \end{pmatrix}.$$

Note that with this parameter set, both monotonicity assumptions are met, so all inference is being done in the case of any assumptions being valid.

5. Generate the 1000 values Y_i from $\text{Bernoulli}(K_i)$ distributions.
6. Analyze the generated Z_i, X_i, Y_i values to form point estimates of the treatment effect in full compliers using each of the following approaches:
- The adapted principal stratification set up in Section 3 using no monotonicity assumptions (using mean of posterior as point estimate). Determine a 95% credible interval using weighted quantiles of the posterior sample.
 - The adapted principal stratification set up in Section 3 using monotonicity assumption 1 (using mean of posterior as point estimate). Determine a 95% credible interval using weighted quantiles of the posterior sample.

- The adapted principal stratification set up in Section 3 using monotonicity assumption 2 (using mean of posterior as point estimate). Determine a 95% credible interval using weighted quantiles of the posterior sample.
- The adapted principal stratification set up in Section 3 using monotonicity assumptions 1 and 2 (using mean of posterior as point estimate). Determine a 95% credible interval using weighted quantiles of the posterior sample.
- Simple principal stratification assuming partial treatment is no treatment. Determine a 95% confidence interval by bootstrapping from the original data set and getting a standard bootstrap confidence interval.
- Simple principal stratification assuming partial treatment is full treatment. Determine a 95% confidence interval by bootstrapping from the original data set and getting a standard bootstrap confidence interval.
- The approach developed in Section 2.2.1 of estimating bounds of the treatment effect in full compliers by assuming monotonicity across principal strata.

After repeating steps 5 and 6 a total of 2000 times, there will then be 2000 sets of 6 point estimates of the effect, as well as 2000 pairs of bound estimates of the effect. Visible in Figure 10 is a comparison of the distributions of each of these sets of 2000 estimates.

Figure 10: Boxplots of estimates of treatment effect in F patients over 2000 simulations using a generative parameter set and population adapted from Shrier[47], separated by estimation methods. The true parameter value is shown as a solid horizontal line.

It was also observed that the true value was within the estimated bounds in 42.3% of simulations.

In addition to examining the performance of single point estimates, it may be observed how coverage of credible/confidence intervals varies across estimation methods, since each point estimation method also had steps included to determine a 95% credible/confidence interval. A plot of 95% interval coverage vs. estimation method is given in Figure 11:

Figure 11: Bar plot of coverage of estimated 95% credible/confidence intervals of treatment effect in F patients over 2000 simulations using a generative parameter set and population adapted from Shrier[47], separated by estimation method.

The main observation here is that in this specific case, assuming partial treatment is equal to receiving no treatment and performing simple principal stratification results in the highest mean square error compared to all other methods, and a coverage below 5%. Also, it is worth observing that the lowest mean square errors are found in any adapted principal stratification method where assumption 1 is made. In comparison, Assumption 2 seems to have little impact on estimation error, but as was seen in Section 3.1.1 and Section 3.1.2, Assumption 1 tends to narrow posterior distributions, so it is not unexpected that Assumption 1's inclusion decreases RMSE (since Assumption 1 is valid) but also decreases coverage. The much lower coverage of the two simple principal stratification methods does not disagree with Shrier's suggestion that these estimation methods are inherently biased, and the higher coverage of the adapted principal stratification approach does not disagree with the suggestion that these methods result in a reduction of this bias.

As a next step in this observation process, we can try observing how these estimator performances change as the sample size increases. If we are correct that the developed estimation methods are unbiased and the simple principal stratification methods are biased, we would expect a lower coverage in the simple estimates' intervals and a high coverage in the new methods' intervals. The above simulations are repeated, but the number of patients in each group is increased by a factor of 8, so there are 8000 total patients instead of 1000. All groupings of patients are also increased by the same factor, so no proportions in the population change. After completing the simulation steps once more with this single change, Figure 12 and Figure 13 are produced:

Figure 12: Boxplots of estimates of treatment effect in F patients over 2000 simulations using a generative parameter set and population adapted from Shrier[47] with total population increased to 8000, separated by estimation method. The true parameter value is shown as a solid horizontal line.

Figure 13: Bar plot of coverage of estimated 95% credible/confidence intervals of treatment effect in F patients over 2000 simulations using a generative parameter set and population adapted from Shrier[47], separated by estimation method, after increasing total sample size to 8000.

It was also now observed that the true value was within the estimated bounds in exactly 52% of simulations. It is visible in these results that as sample size increased, all estimates have reduced variance. We also see that the coverage of both simple principal stratification 95% intervals now approach zero, further agreeing that these estimation methods are inherently biased.

We may also see in Table 2 the observed mean square error resulting from each of these estimation methods (except the bound estimation), and see how these measures of error changed as sample size increased.

Table 2: Mean square errors of six principal stratification (PS) estimation methods from 2000 simulations completed using a generative parameter set and population adapted from Shrier[47], in the case with a sample size of 1000 and in an altered case with a total population of 8000.

Estimation Method	Root Mean Square Error	
	Sample Size 1000	Sample Size 8000
Adapted PS, No Assumption	0.0781	0.0689
Adapted PS, Assumption 1	0.0565	0.0367
Adapted PS, Assumption 2	0.0975	0.0867
Adapted PS, Assumptions 1 & 2	0.0506	0.0350
Simple PS, Assume Partial as None	0.273	0.268
Simple PS, Assume Partial as Full	0.0878	0.0794

It is visible that while the posterior mean estimator RMSE values all reduced to less than 90% of their original value, the simple principal stratification estimators assuming partial treatment as no treatment or as full treatment only had RMSE reduce to 98% and 91% of their original values respectively. In both cases the best performance was seen in the posterior mean of the adapted principal stratification, under any assumption setup except assumption 2 without assumption 1. These findings do agree with the simulations shown in Section 3.1.1 and 3.1.2, which indicated that assumption 1 does tend to narrow the posterior, while assumption 2 can result in more irregular shape in the posterior distribution which is even more exaggerated when assumption 2 is made without assumption 1. For more theory on coverage of credible intervals in partially identified models please see Gustafson [20].

4 Conclusions, Future Work

Bias was identified in standard clinical trial inference methods when faced with partial adherence to treatment protocol. Also described was the lack of more advanced methods that are feasible in patient-oriented research settings. New methods based on principal stratification were presented and have been shown to avoid the biases of previous approaches. One of the newly proposed methods is a simple bound calculation and one is a posterior sampling

method. The simple bound calculation method is seen to be consistent as long as partial adherers can be accurately assumed to have outcomes between those of non-adherer and full-adherers (described as Assumption 1), which may be a concern in settings where the treatment period is long enough for adherence to be impacted by treatment outcome. The posterior sampling method is more versatile and can be applied in situations where Assumption 1 is not valid, and can also explore the optional second assumption that treatment effect is monotone with the amount of treatment received (described as Assumption 2). Even in cases where Assumption 1 was valid, the coverage of 95% credible intervals calculated from posterior samples did outperform the coverage of estimated bounds based on the simple bound calculation method, although this was not explored across a full range of parameter sets. The evidence shown here does suggest that the posterior sampling scheme is superior to the simple bound estimation in all regards except for simplicity.

A range of simulations were completed which explored the impact of incorrectly making either monotonicity assumption while sampling posterior estimates. Also observed was the benefit seen when correctly making either assumption. Overall, the impact of Assumptions 1 and 2 seemed varied enough that it proved challenging to try to describe the general expected outcome if either of them was made correctly or incorrectly. One key observation is that Assumption 1 tended to narrow posterior distribution ranges, while Assumption 2 did not, and only altered the shape of the distribution in a similar range. Assumption 2 seldom actually improved estimation, while Assumption 1 often improved it, even in low-sample-size cases where Assumption 1 was not valid, since the bias was small enough to still allow the correct parameter values to fall within the resulting narrowed credible interval. In a high-sample-size case, it would be reasonable to worry that incorrectly making Assumption 1 would result in a bias that outweighed the benefit of a narrowed posterior distribution, but this was also not fully explored.

The findings presented here are built on previously developed theory on principal stratification [42], pragmatic clinical trials [38, 40, 14, 45, 10, 50], and partially identified models [21, 20]. They add to and do not contradict the findings of Shrier et. al. [47] that reported bias in cases where partial treatment was recorded as full treatment or no treatment in clinical trials. The findings of this thesis took a different approach than some previously presented methods that are less feasible in pragmatic clinical research settings, but do remain useful in settings where a larger amount of information can be collected from patients [17, 32, 39, 31, 36, 5, 29, 18].

This research is useful to caregivers, because it highlights that analytical tools exist for reducing bias in settings of partial treatment adherence that don't require incredibly strict regimens of collecting patient information. Using the presented methods, patients could be enrolled in clinical research with minimal extra effort, which would be of great benefit given the unavoidable added stress faced when many steps are necessary to access treatment. The methods presented in this thesis can already be used for great benefit in binary-outcome settings where patients can not access treatment unless granted access. However, these methods could also potentially be expanded in future research to more complex settings using the already presented framework. The most evident next steps to explore the use of these developed methods are to consider expanding them to cases of continuous responses or more diverse principal strata. While the methods of using partially identified model theory to create a more straightforward Monte Carlo sampling method were unique to the binary outcome and simple three-strata framework, a slower Markov Chain Monte Carlo method could be used to form estimates in more complicated strata setups or for continuous response data.

Overall, these methods have a promise that suggests there would be a great benefit in clinical trials moving forward to keep track of when treatments are partially administered. Current methods of assigning a threshold amount of treatment to count as receiving the full

treatment should be abandoned for methods that properly acknowledge when treatment has been administered but not in full. Hopefully, as a result, biased results can be avoided in future pragmatic clinical trial research.

References

- [1] J. Angrist, G. Imbens, and D. Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91:444–455, Jun 1993.
- [2] P. C. Austin. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate Behavioral Research*, 46(3):399–424, 2011.
- [3] A. Balke and J. Pearl. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association*, 92(439):1171–1176, 1997.
- [4] J. H. Baron. Evolution of clinical research: A history before and beyond james lind. *Perspectives in Clinical Research*, 3(4):149, 2012.
- [5] F. Bartolucci and L. Grilli. Modeling partial compliance through copulas in a principal stratification framework. *Journal of the American Statistical Association*, 106(494):469–479, 2011.
- [6] R. Campbell, F. Starkey, J. Holliday, S. Audrey, M. Bloor, N. Parry-Langdon, R. Hughes, and L. Moore. An informal school-based peer-led intervention for smoking prevention in adolescence (assist): a cluster randomised trial. *The Lancet*, 371(9624):1595–1602, 2008.
- [7] J. Cheng and D. S. Small. Bounds on causal effects in three-arm trials with non-compliance. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(5):815–836, 2006.
- [8] N. K. Choudhry, J. Avorn, R. Glynn, E. Antman, S. Schneeweiss, M. Toscano, L. Reisman, J. Fernandes, C. Spettell, J. L. Lee, and et al. Full coverage for preventive medications after myocardial infarction. *New England Journal of Medicine*, 365:2088–2097, 2011.

- [9] Coronary Drug Project Research Group. Influence of adherence to treatment and response of cholesterol on mortality in the coronary drug project. *Journal of Occupational and Environmental Medicine*, 23(7):507, 1981.
- [10] S. Dodd, I. R. White, and P. Williamson. Nonadherence to treatment protocol in published randomised controlled trials: a review. *Trials*, 13(1), 2012.
- [11] I. Ford and J. Norrie. Pragmatic trials. *New England Journal of Medicine*, 375(5):454–463, Apr 2016.
- [12] C. E. Frangakis, D. B. Rubin, and X. H. Zhou. Clustered encouragement designs with individual noncompliance: bayesian inference with randomization, and application to advance directive forms. *Biostatistics*, 3(2):147–164, 2002.
- [13] O. Fröbert, B. Lagerqvist, G. K. Olivecrona, E. Omerovic, T. Gudnason, M. Maeng, M. Aasa, O. Angerås, F. Calais, M. Danielewicz, and et al. Thrombus aspiration during st-segment elevation myocardial infarction. *New England Journal of Medicine*, 369(17):1587–1597, 2013.
- [14] Government of Canada Canadian Institutes of Health Research. Strategy for patient-oriented research. *CIHR*, Jan 2019.
- [15] S. Greenland. An introduction to instrumental variables for epidemiologists. *International Journal of Epidemiology*, 47(1):358–358, 2017.
- [16] R. Greevy, J. H. Silber, A. Cnaan, and P. R. Rosenbaum. Randomization inference with imperfect compliance in the ace-inhibitor after anthracycline randomized trial. *Journal of the American Statistical Association*, 99(465):7–15, 2004.
- [17] S. F. Grey. *How Much Compliance is Enough? Examining the Effect of Different Definitions of Compliance on Estimates of Treatment Efficacy in Randomized Controlled Trials*. PhD thesis, Case Western Reserve University, 2013.

- [18] T. Guo. *Causal effects in randomized trials in the presence of partial compliance: breast-feeding effect on infant growth*. PhD thesis, McGill University, 2009.
- [19] S. Gupta. Intention-to-treat concept: A review. *Perspectives in Clinical Research*, 2(3):109, 2011.
- [20] P. Gustafson. On the behaviour of bayesian credible intervals in partially identified models. *Electronic Journal of Statistics*, 6:2107–2124, 2012.
- [21] P. Gustafson. *Bayesian inference for partially identified models: exploring the limits of limited data*. CRC Press, Boca Raton, 2015.
- [22] E. F. Halpern. Behind the numbers: Inverse probability weighting. *Radiology*, 271(3):625–628, 2014.
- [23] Heart Protection Study Collaborative Group. Mrc/bhf heart protection study of cholesterol lowering with simvastatin in 20,536 high-risk individuals: a randomised placebo-controlled trial. *The Lancet*, 360:7–22, 2002.
- [24] M. A. Hernán, S. Hernández-Díaz, and J. M. Robins. Randomized trials analyzed as observational studies. *Annals of Internal Medicine*, Oct 2013.
- [25] M. A. Hernán and S. Hernández-Díaz. Beyond the intention-to-treat in comparative effectiveness research. *Clinical Trials: Journal of the Society for Clinical Trials*, 9(1):48–55, 2011.
- [26] M. A. Hernán and J. M. Robins. Per-protocol analyses of pragmatic trials. *New England Journal of Medicine*, 377(14):1391–1398, May 2017.
- [27] M. A. Hernán and D. Scharfstein. Cautions as regulators move to end exclusive reliance on intention to treat. *Annals of Internal Medicine*, 168(7):515, 2018.

- [28] W. Huiping, W. Yu, J. Pei, L. Jiao, Z. Shian, J. Hugang, W. Zheng, and L. Yingdong. Compound salvia pellet might be more effective and safer for chronic stable angina pectoris compared with nitrates. *Medicine*, 98(9), Mar 2019.
- [29] G. W. Imbens and D. B. Rubin. Bayesian inference for causal effects in randomized experiments with noncompliance. *The Annals of Statistics*, 25(1):305–327, 1997.
- [30] G. W. Imbens and D. B. Rubin. Bayesian inference for causal effects in randomized experiments with noncompliance. *The Annals of Statistics*, 25(1):305–327, 1997.
- [31] H. Jin and D. B. Rubin. Principal stratification for causal inference with extended partial compliance. *Journal of the American Statistical Association*, 103(481):101–111, 2008.
- [32] H. Jin and D. B. Rubin. Public schools versus private schools: Causal inference with partial compliance. *Journal of Educational and Behavioral Statistics*, 34(1):24–45, 2009.
- [33] L. Kish. *Survey sampling*. Wiley-Interscience publication, New York, 1965.
- [34] P. M. Lee. *Bayesian statistics: an introduction*. Wiley, Chichester, 2012.
- [35] R. J. Little, Q. Long, and X. Lin. A comparison of methods for estimating the causal effect of a treatment in randomized clinical trials subject to noncompliance. *Biometrics*, 65(2):640–649, 2008.
- [36] Y. Ma and J. Roy. Causal models for randomized trials with continuous compliance. *Statistical Causal Inferences and Their Applications in Public Health Research ICSCA Book Series in Statistics*, pages 187–201, 2016.
- [37] E. J. Murray and M. A. Hernán. Adherence adjustment in the coronary drug project: A call for better per-protocol effect estimates in randomized trials. *Clinical Trials: Journal of the Society for Clinical Trials*, 13(4):372–378, Jul 2016.

- [38] National Institute of Health Research. About involve. *INVOLVE*.
- [39] L. C. Page. Understanding the impact of career academy attendance. *Evaluation Review*, 36(2):99–132, 2012.
- [40] Patient-Centered Outcomes Research Institute. Pcori. *Dissemination and Implementation Framework and Toolkit*, Jan 2019.
- [41] N. A. Patsopoulos. A pragmatic view on pragmatic trials. *Dialogues in Clinical Neuroscience*, 13(2):217, 2011.
- [42] J. Pearl. Principal stratification – a goal or a tool? *The International Journal of Biostatistics*, 7(1):1–13, 2011.
- [43] R. Pickard, T. Lam, G. MacLennan, K. Starr, M. Kilonzo, G. McPherson, and et al. Antimicrobial catheters for reduction of symptomatic urinary tract infection in adults requiring short-term catheterisation in hospital: a multicentre randomised controlled trial. *The Lancet*, 380(9857):1927–1935, Nov 2012.
- [44] I. Roberts, D. Yates, P. Sandercock, B. Farrell, J. Wasserberg, G. Lomas, R. Cottingham, P. Svoboda, N. Brayley, G. Mazairac, and et al. Effect of intravenous corticosteroids on death within 14 days in 10008 adults with clinically significant head injury (mrc crash trial): randomised placebo-controlled trial. *Obstetrics & Gynecology*, 105(1):214, 2005.
- [45] D. Schwartz and J. Lellouch. Explanatory and pragmatic attitudes in therapeutical trials. *Journal of Clinical Epidemiology*, 62(5):499–505, 2009.
- [46] A. S. V. Shah, A. Anand, F. E. Strachan, A. V. Ferry, K. K. Lee, A. R. Chapman, and et al. High-sensitivity troponin in the evaluation of patients with suspected acute coronary syndrome: a stepped-wedge, cluster-randomised controlled trial. *The Lancet*, 392:919–928, 2018.

- [47] I. Shrier, R. W. Platt, R. J. Steele, and M. Schnitzer. Estimating causal effects of treatment in a randomized trial when some participants only partially adhere. *Epidemiology*, 29(1):78–86, 2018.
- [48] I. Shrier, E. Verhagen, and S. D. Stovitz. The intention-to-treat analysis is not always the conservative approach. *The American Journal of Medicine*, 130(7):867–871, 2017.
- [49] K. E. Thorpe, M. Zwarenstein, A. D. Oxman, S. Treweek, C. D. Furberg, D. G. Altman, S. Tunis, E. Bergel, I. Harvey, D. J. Magid, and et al. A pragmatic-explanatory continuum indicator summary (precis): a tool to help trial designers. *Canadian Medical Association Journal*, 180(10), 2009.
- [50] T. J. Vanderweele. Mediation analysis with multiple versions of the mediator. *Epidemiology*, 23(3):454–463, 2012.