Improving Bayesian adaptive clinical trials with covariate information

James Willard

Doctor of Philosophy

Department of Epidemiology, Biostatistics and Occupational Health

McGill University
Montréal, Québec
May 2024

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Doctor of Philosophy

Dedication

For my mother.

"May your choices reflect your hopes, not your fears."

– Nelson Mandela

Acknowledgements

There are many people I would like to thank for making this thesis possible. Thank you to my mother Carol, stepfather Steve, and sister Morgan for their incredible love and support throughout my life. To Pop Pop and my grandmother Elaine, thank you for impressing upon me the importance of thinking critically and for teaching me the value in doing the hard things. To my grandmother Lois, thank you for always believing in me, for teaching me to think creatively, and for your counsel and constant encouragement every step of the way.

I am lucky to have such amazing friends. Thank you to Taylor, Raleigh, and Nicole, for cheering me on and keeping me going all of these years. Thank you to Gabrielle, for introducing me to all of the wonderful things Montreal has to offer, for helping me practice my French, and for providing me a family away from home. To Eduardo, thank you for your friendship, for all of the wonderful dinners and parties, and for the positive encouragement you gave me throughout the PhD. To the many friends I have made at McGill, including Marc, Vanessa, Niki, Julien, and Victoire, thank you for the cinq à septs, the really fun nights out, and for making the department a warm and welcoming place to work and study.

I am grateful for the education I received while at McGill, which is of course a reflection of the wonderful faculty who instructed me. Thank you to Jim Hanley, Andrea Benedetti, and Michal Abrahamowicz for their courses which I greatly enjoyed. To Alex Schmidt, thank you for teaching me, for providing me a chance to attend ISBA in Montreal, and for being a passionate advocate of Bayesian statistics. I am grateful for the experiences I had while I was an intern in industry. Thank you to Brad Carlin for involving me in a variety of engaging projects and for taking an interest in me and my research. To Bruno Boulanger, thank you for supporting my many conference presentations; they greatly enriched my graduate

student experience. Finally, thank you to Lira Pi for being a great person to work with.

And last, but certainly not least, thank you to my amazing supervisors, Shirin Golchi and Erica Moodie, without whom none of this would have been possible. To Shirin, thank you for tailoring my experience as a student to my individual goals and interests. I really appreciate you finding and supporting many of the opportunities I took advantage of, including paper competitions, conference presentations, and research exchanges. I am grateful for having had the opportunity to work with you on projects outside of my thesis; I definitely learned a lot. To Erica, thank you for always taking the time to clarify concepts for me and for your detailed comments on my work. Your thoughtful input has helped me improve my writing, in both structure and style. I really appreciate how you were always a great resource to turn to, and for your positive encouragement along the way, which was especially helpful and reassuring during some of the more difficult times of the PhD.

I am extremely grateful for Shirin and Erica having mentored me over the past four years and for them patiently guiding me and challenging me in ways that helped me grow as a person and thinker. I am a better statistician because of them. Any strengths of this thesis are a direct result of their steadfast commitment to my education and development as a researcher. All errors and omissions are my own.

Preface

This manuscript-based thesis is comprised of six chapters: an introduction, a comprehensive literature review, three chapters corresponding to individual manuscripts, and a conclusion. At the end of the thesis, a single appendix is presented which is followed by a complete bibliography. Chapters 3, 4, and 5 provide methodological contributions to both the late phase randomized controlled trial literature and the early phase dose-finding trial literature. Furthermore, they are united by the central theme of this thesis, which is concerned with incorporating covariate information into Bayesian adaptive clinical trials. Each of these chapters begins with a preamble that briefly introduces the context of the manuscript and identifies the knowledge gap it seeks to fill. Each manuscript includes a real-world example which demonstrates a realistic application of its corresponding proposed method.

The introduction and literature review (Chapters 1 and 2) were conceived and written by James Willard (JW). These chapters were improved through revision following an insightful discussion with, and helpful comments from, Shirin Golchi (SG) and Erica E.M. Moodie (EEMM). The conceptualization of the work contained in Chapter 3 resulted from a series of discussions with SG and EEMM. JW was responsible for deriving the methods, designing and implementing the simulation study and application, and writing the draft manuscript. SG and EEMM provided substantial guidance and assistance in deriving the methods and additionally helped verify and interpret the results obtained from the simulations. SG and EEMM also provided helpful edits for the chapter. The methods proposed in Chapter 4 were conceptualized by JW, SG, EEMM, and Brad Carlin (BC). They were defined as part of JW's internship at a host partner from industry, where BC and Bruno Boulanger (BB) were the industry contacts. JW derived the methods, designed and performed the simulation study and application, and wrote the draft manuscript. BB provided valuable input regarding the simulation study. The work was advised and edited by SG, EEMM, and BC. The ideas in

Chapter 5 were conceptualized by JW, SG, and EEMM. JW derived the methods, designed and performed the simulation study and application, and wrote the draft manuscript. SG and EEMM advised and edited the work in the chapter. The conclusion of this thesis was conceived and written by JW and was edited by SG and EEMM.

Abstract

Bayesian adaptive clinical trials allow for predetermined changes to the trial design based on evidence provided by the accumulating data. Early phase trials assess the safety and efficacy profiles of first-in-human doses of an experimental intervention, whereas late phase trials seek to confirm the efficacy of an experimental intervention. It is common in both phases of trials to collect information on many covariates, though less common to utilize this information within the trials. Through three manuscripts, this thesis investigates the benefit of using this additional information in both early phase dose-finding trials for combination therapies and late phase confirmatory trials. In the first manuscript, I focus on late phase confirmatory trials, where I propose adjusting for covariates known to be at least moderately associated with the outcome. Adjusted and unadjusted analyses are described within the context of collapsibility and non-collapsibility of commonly used estimands in clinical trials, and a tutorial for obtaining marginal treatment effect estimates from adjusted analyses is provided. It is shown that covariate adjustment increases the power and probability of stopping trials early, and decreases expected sample sizes, demonstrating covariate adjustment leads to trials which stop earlier and more often. This is followed by a second manuscript which proposes using Bayesian optimization methods for early phase personalized dose-finding trials with combination therapies, where individual patient characteristics are used to recommend patient-specific optimal dose combinations. Focus is placed on a minimal toxicity setting where formal dose escalation rules are not required. The approach demonstrates the feasibility of a personalized approach for dose-finding trials despite the limited sample sizes which are common in early phase trials. The third and final manuscript generalizes this approach to the setting of higher-grade toxicities, where a formal dose escalation scheme is introduced and where toxicities are incorporated into the personalized dose-finding strategy. In summary, this thesis showcases the benefit of utilizing additional covariate information in the context of Bayesian adaptive clinical trials and seeks to promote more widespread use of the methods proposed herein.

Abrégé

Les essais cliniques adaptatifs bayésiens permettent d'apporter des modifications prédéterminées à leurs conceptions sur la base des preuves fournies par les données qui s'y accumulent. Les essais de phase précoce évaluent les profils de sécurité et d'efficacité des premières doses d'une intervention expérimentale administrées à l'homme. Les essais de phase tardive cherchent quant à eux à confirmer l'efficacité d'une intervention expérimentale. Dans les deux phases d'essais, il est courant de recueillir des informations sur de nombreuses variables, mais il est moins courant d'utiliser ces informations dans le cadre de ceux-ci. Au travers de trois manuscrits, cette thèse étudie les avantages associés à l'utilisation de ces informations supplémentaires à la fois dans la recherche de dose pour les thérapies combinées en phase précoce que dans les essais de confirmation en phase tardive. Le premier manuscrit porte sur les essais de confirmation de phase tardive, où il est proposé de prendre en compte des variables connues comme étant au moins modérément associées au résultat. Des analyses ajustées et non ajustées sont décrites dans le contexte de la collapsibilité et de la non-collapsibilité des estimateurs couramment utilisés dans les essais cliniques, et un tutoriel pour l'obtention d'estimations de l'effet marginal du traitement à partir d'analyses ajustées est fourni. Il y est démontré que l'ajustement augmente la puissance et la probabilité d'arrêter les essais précocement et diminue les tailles d'échantillon attendues, démontrant que l'ajustement conduit à des essais qui s'arrêtent plus tôt et plus souvent. Il est suivi d'un second manuscrit qui propose d'utiliser des méthodes d'optimisation bayésienne lors de la phase précoce de recherche de doses personnalisées pour des thérapies combinées, où les caractéristiques individuelles du patient sont utilisées pour recommander des combinaisons de doses optimales spécifiques au patient. L'accent y est mis sur un cadre de toxicité minimale où les règles formelles d'escalade des doses ne sont pas nécessaires. Cette méthode démontre la faisabilité d'une approche personnalisée lors d'essais de détermination de la dose, malgré la taille limitée des échantillons fréquente dans les essais de première phase. Le troisième et dernier manuscrit généralise cette approche dans le cadre de toxicités plus importantes, où

un schéma formel d'escalade de dose est introduit et où des toxicités sont incorporées dans la stratégie personnalisée de recherche de dose. En résumé, cette thèse met en évidence les avantages de l'utilisation d'informations supplémentaires dans le contexte d'essais cliniques adaptatifs bayésiens et cherche à promouvoir une utilisation plus répandue des méthodes proposées ici.

Table of contents

List of Tables

List of Figures

Abbreviations

AE adverse event

FDA United States Food and Drug Administration

GP Gaussian process

MTD maximum tolerated dose

OBD optimal biological dose

RCT randomized controlled trial

Chapter 1

Introduction

This thesis is concerned with the incorporation of covariate information into Bayesian adaptive clinical trials. I focus on both late phase randomized controlled trials (RCTs) and early phase dose-finding trials, and consider how utilizing covariate information might be beneficial within each.

In any clinical trial, it is common to collect patient information which is not directly related to treatment assignment or response. This thesis argues that utilizing this additional covariate information within the scope of the trial is beneficial. In conventional, fixed size RCTs, adjustment for variables known to be associated with the outcome has been shown to increase statistical power (e.g., Benkeser et al. (2021)). These power increases have recently been shown in more flexible frequentist designs, such as group sequential and information adaptive designs (Van Lancker et al., 2022) as well as for adaptive multi-arm designs (Lee et al., 2022). Despite this benefit, covariate adjustment has not been characterized within flexible Bayesian designs which include adaptive stopping rules. Furthermore, the impact of combining prior information with covariate adjustment has not been previously explored.

In early phase dose-finding trials, the additional covariate information may be used to tai-

lor doses to individual patients. This *personalized* approach to dose-finding recognizes that patient response heterogeneity may exist in the population, which stands in stark contrast to *standard* dose-finding methods which assume a single dose applies equally well to all patients. Extending many standard parametric dose-finding approaches to the personalized setting is challenging, due to the limited sample sizes but potentially large number of treatment-covariate interaction terms which must be estimated. These estimation challenges are exacerbated in the combination therapy setting, where two or more dosing agents are combined. Mozgunov et al. (2022) has considered dual-agent personalized dose-finding, but where the patient-specific dose of one of the agents is selected externally by clinicians. The literature for personalized dose-finding trials where tailoring is performed with respect to multiple dosing agents remains underdeveloped.

In this thesis, I build upon the previous statistical literature investigating covariate adjustment in late phase RCTs, and show how it improves the operating characteristics of Bayesian designs which include adaptive stopping rules. I also propose adapting methods from the Bayesian optimization literature to the setting of personalized dose-finding trials with combination therapies, which permits multiple dosing agents to be tailored to individual patients despite the small sample sizes in these trials. In Chapter 2, I review the most important concepts which are required to understand the methods that are developed in the chapters that follow. Specifically, I discuss conducting adaptive clinical trials under the Bayesian paradigm, and describe key features of both late phase RCTs and early phase personalized dose-finding trials.

In Chapter 3, I describe covariate adjustment in the context of late phase Bayesian adaptive RCTs. I discuss the collapsibility and non-collapsibility of estimands for continuous, binary, and time-to-event endpoints. As is common in RCTs, I focus on marginal estimands and provide a tutorial for obtaining marginal posterior estimates from adjusted analyses. Through simulation, I compare the impact of covariate adjustment and the incorporation of different

strengths of prior information for the covariate effects on the operating characteristics of trials with different maximum sample sizes and endpoints. I then demonstrate the benefit of covariate adjustment in a design which investigates the effectiveness of oral therapies against mild to moderate COVID-19 infection in individuals discharged from Canadian Emergency Departments.

In Chapter 4, I consider early phase Bayesian adaptive personalized dose-finding trials for combination therapies under a minimal toxicity setting. The sequential dose-finding algorithm proposed utilizes methods from the Bayesian optimization literature. Specifically, I propose estimating the dose-efficacy surface using a Gaussian Process (GP) model, and then define a sequential search policy using an acquisition function which is conditioned on patient covariate patterns. Through simulation, I investigate and contrast the performance of both the personalized and standard dose-finding approaches under scenarios which either include or exclude patient response heterogeneity. I then compare a set of proposed designs for an industry-sponsored problem which is concerned with a design for the development of an intraocular implant that combines two topical agents.

In Chapter 5, I generalize the personalized dose-finding approach of Chapter 4 to the setting of higher-grade toxicities. Under this scenario, toxicity information must be formally incorporated into the dose-finding. To do so, I assume the dose-efficacy and dose-toxicity surfaces are conditionally independent given dose, and estimate them using independent GP models. I define a sequential search policy through an acquisition function which includes patient-specific toxicity constraints. Furthermore, I propose a dose-escalation scheme to collect initial patient responses in the presence of potentially toxic doses. Through simulation, I assess the performance of both the personalized and standard approaches in the presence and absence of patient response heterogeneity, and compare a set of proposed designs for the development of a dual-agent pharmacotherapy for obstructive sleep apnea.

Chapters 3 to 5 were written as individual manuscripts. Chapter 3 is published in *Statistical*

Methods in Medical Research. Chapters 4 and 5 are currently under review in statistical journals. An effort has been made to preserve the notation across chapters. I end this thesis with a conclusion, which reviews the methodological contributions made by the thesis, discusses important limitations of the proposed methods, and describes extensions and possible new directions for future work.

Chapter 2

Literature review

In this chapter, I discuss the core ideas and concepts which are required to understand the content of this thesis. First, I review clinical trials and illustrate some of the challenges associated with modern day drug development. I then introduce Bayesian adaptive clinical trials, a flexible class of trials on which this thesis focuses, and in which the posterior distribution is the primary quantity of interest. I review the adaptive stopping rules used in these trials and discuss how they satisfy the likelihood principle. In a third section, I focus on late phase RCTs where I review how covariate adjustment has been used to improve trial operating characteristics and to permit a personalized approach toward medicine. After this, I shift focus to early phase designs and describe historical approaches toward dose-finding and the challenges of evaluating combination therapies in a personalized medicine setting. A final section discusses Bayesian optimization within the context of early phase dose-finding.

2.1 Clinical trials

Clinical trials help researchers address scientific questions about an experimental intervention (hereafter referred to as a "treatment", "therapy", or "drug") in human populations. Clinical experimentation in human beings has a long history, with Bull (1959) tracing some

of the earliest recorded examples back to the medical practices of the ancient Egyptians around 2,000 B.C.E. Many core principles now commonly used in clinical trials, e.g., use of a control group and randomization, were only introduced within the past few hundred years. For example, one of the earliest clinical trials which contained an untreated control group is from 1747, where the surgeon James Lind compared several treatments for scurvy, a disease that is caused by severe vitamin C deficiency and that was highly prevalent among sailors at the time (Lind, 1753). While this trial did contain a control group, it was not randomized. Randomization would be introduced to experimental design by Ronald Fisher in 1923, through his work in agricultural field experiments (Fisher and Mackenzie, 1923). Confirmatory clinical trials would eventually adopt these methods, with one of the earliest examples of an RCT being performed by the Medical Research Council of the United Kingdom in 1948, where streptomycin was investigated as a potential treatment for pulmonary tuberculosis, a potentially fatal infection that primarily affects the lungs (Medical Research Council, 1948). One member of the trial's committee, Sir Austin Bradford Hill, would become a major leader in the development of modern day clinical trials, eventually producing a highly influential text in the field, *Statistical Methods in Clinical and Preventive Medicine* (Bradford Hill, 1962).

Over time, the continued development of clinical trials have led them to become considered as the "gold standard" method for evaluating experimental interventions. Modern day clinical trials are typically performed through either the frequentist (e.g., Pocock (1983)) or Bayesian (e.g., Spiegelhalter et al. (2004)) paradigm of statistics, though in this thesis, I focus specifically on trials which utilize the latter. Adopting a Bayesian approach in clinical trials is advantageous since it naturally accommodates the sequential nature of adaptive designs, it provides valid statistical inference even in small sample sizes (Bernardo and Smith, 1994), and it facilitates the incorporation of prior information and/or historical controls into a trial's analysis (e.g., Ibrahim and Chen (2000)). Prior information may come in the form of clinical experience, previous animal or clinical studies, and possibly from biology in the

form of pharmacokinetic/pharmacodynamic information (e.g., Gelman et al. (1996)). The United States Food and Drug Administration (FDA) often categorizes clinical trials into four distinct phases, each of which has different objectives, is typically of a different size, and has different success rates as described in the following. Trials in Phase 1 seek to discover treatment dosages that are well tolerated (i.e., safe) and that may suggest some benefit for further study. These trials are typically not randomized and enroll roughly 20-100 participants, with approximately 70% of interventions moving on to Phase 2 (United States Food and Drug Administration, 2018b). Trials in Phase 2 continue to collect safety information and investigate whether the experimental treatment suggests efficacy, though do not include enough participants to determine this definitively. These trials may or may not be randomized and enroll up to several hundred participants, with roughly 33% of interventions moving on to Phase 3 (United States Food and Drug Administration, 2018b). Trials in Phase 3 are typically randomized and serve to confirm the efficacy of an experimental treatment. These trials also permit a greater assessment of a treatment's safety, as they typically range from several hundred to tens of thousands of participants depending on the type of intervention and outcome of interest, and often last for several years or more. Approximately 25-30% of treatments are approved at this stage and may be monitored over time in Phase 4 Post-Market Safety Monitoring studies (United States Food and Drug Administration, 2018b).

Requiring experimental treatments to undergo evaluation through the phases described above makes drug development an expensive, time-consuming process (Avorn, 2015). As trials are additionally plagued by high failure rates, there is growing interest in making this process more efficient and successful overall (Kimko and Pinheiro, 2015). One proposed solution toward this end is increasing the utilization of more flexible adaptive designs, which can make trials more efficient while preserving the (frequentist) trial operating characteristics commonly required by regulators (e.g., Type 1 error rate and statistical power; Pinheiro et al. (2011)). Contrasted with fixed Bayesian designs where the trial plan and sample size

are specified in advance (Chaloner and Verdinelli, 1995), Bayesian adaptive designs permit pre-defined trial changes based upon the accumulating data (Berry et al., 2010). In this thesis, I specifically focus on Bayesian adaptive designs.

2.2 Bayesian adaptive clinical trials

Bayesian adaptive clinical trials are sequential in nature and utilize the accumulating data to entertain a variety of adaptations, which follow pre-specified rules and may include adaptive allocation, adaptive sampling, adaptive stopping, and adaptive enrichment (Giovagnoli, 2021). Adaptive allocation rules change the randomization procedure over the course of the trial. For example, response adaptive randomization updates the allocation probability to a particular treatment arm based on previously observed responses, with the goal of assigning more participants to better performing treatments (e.g., Berry and Eick (1995); Thall and Wathen (2007)). Adaptive sampling rules may permit sample size re-estimation at an interim analysis which can be useful for trials with long accrual periods and for which great uncertainty exists regarding the parameter estimates used for trial design and planning (e.g., Uemura et al. (2017)). Adaptive stopping rules, to be reviewed in detail below, permit the trial to stop early for any combination of safety, futility, or efficacy of treatment. Adaptive enrichment designs may permit changes to eligibility criteria, add or drop treatments, or assess the effectiveness of therapies within subgroups of the population that appear more likely to benefit from them (see Thall (2021)).

To further improve efficiency, so-called "seamless" Bayesian adaptive designs have been proposed which combine participant data across, and permit transitioning between, several trial phases in a single, unified study. Seamless designs have been proposed to combine Phases 1 and 2 (see Yuan et al. (2016b)), Phases 2 and 3 (see Thall (2008)), and even Phases 1, 2, and 3 (e.g., Chapple and Thall (2019)). Unless explicitly required, I do not focus on the specific distinctions between trial phases, but will instead use *early phase* to broadly refer to Phases 1 or 2 and *late phase* to refer to Phase 3. I do not consider Phase 4 clinical trials

in this thesis. Recently, so-called Bayesian adaptive "master protocol" designs have been proposed to assess one or more targeted therapies in one or more diseases or subtypes under a single clinical trial protocol, and are generally divided into "umbrella", "basket", or "platform" trials (Woodcock and LaVange, 2017). Trials of these types may facilitate advances in the field of personalized medicine, to be reviewed in detail in the next section. A prime example of the flexibility of adaptive designs is the I-SPY2 platform trial, which investigates targeted therapies for early-stage breast cancer and is still ongoing at the time of writing (Barker et al., 2009). The I-SPY2 protocol adds new treatments as they become available, adaptively stops enrollment in treatments which are deemed to be ineffective, adaptively randomizes participants to better performing treatment arms based on their tumor subtype, and "graduates" (i.e., promotes to a later phase of study) treatments which are deemed likely to be beneficial.

2.2.1 Posterior and posterior predictive distributions

The decision to perform a specific adaptation is governed by predefined decision criteria, which are commonly defined as posterior or posterior predictive probability statements about a parameter of interest. The posterior distribution is a central focus in Bayesian statistics and represents one's uncertain beliefs about unknown parameters of interest, $\boldsymbol{\theta}$, in light of all data which have been previously observed, $\mathbf{y}$. To obtain the posterior distribution, we represent our initial beliefs about $\boldsymbol{\theta}$ in the form of a prior distribution $p(\boldsymbol{\theta})$. These initial beliefs are then updated by data through the likelihood function, $p(\mathbf{y} \mid \boldsymbol{\theta})$. The updated beliefs are represented by the posterior distribution, $p(\boldsymbol{\theta} \mid \mathbf{y})$, which is obtained using Bayes' Theorem (Gelman et al., 2014):

$$p(\boldsymbol{\theta} \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid \boldsymbol{\theta})p(\boldsymbol{\theta})}{\int_{\boldsymbol{\theta}} p(\mathbf{y} \mid \boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta}}. \tag{2.1}$$

The denominator to the right of the equality serves as a normalizing constant to ensure the posterior is a properly scaled distribution. The posterior predictive distribution represents

one's uncertain beliefs about the value of a future observation y^* given the previously observed data. This distribution, denoted by $p(y^* \mid \mathbf{y})$, fully accounts for the uncertainty in $\boldsymbol{\theta}$ when predicting y^* by marginalizing over the posterior distribution of $\boldsymbol{\theta}$:

$$p(y^* \mid \mathbf{y}) = \int p(y^* \mid \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid \mathbf{y}) d\boldsymbol{\theta}. \tag{2.2}$$

Since the decision criteria determining whether or not to execute a specific adaptation are typically functionals of these distributions, the performance of a specific design often depends on how well they are estimated. When these distributions are not known analytically, they must be approximated, with one of the most commonly used approximation methods being Markov chain Monte Carlo (e.g., Gelfand and Smith (1990)).

2.2.2 The likelihood principle

Early proponents of the Bayesian paradigm advocated for its use in sequential trials because it is consistent with the *likelihood principle* (e.g., (Cornfield, 1966)). I denote the likelihood by $L_{\mathbf{y}}(\theta)$, viewed as a function of unknown parameter θ, with $\mathbf{y}$ representing the observed data. The likelihood principle states that $L_{\mathbf{y}}(\theta)$ contains all information from an experiment that is relevant for performing inference on θ, and that two likelihood functions which differ only by a scalar multiple should yield the same inference (Berger and Wolpert, 1988). This means that the same conclusions should be reached from two samples which yield proportional likelihood functions. If one accepts the likelihood principle, then inference should not depend on the stopping rule used to collect the data. Bayesian adaptive clinical trials which employ early stopping rules are consistent with this principle. To illustrate this idea, I consider the following example adapted from Berry (1987).

Suppose there is a clinical trial which assigns two treatments, A and B, to each participant. After undergoing both treatments, each participant selects a preferred one, where ties are not permitted. Let θ denote the population proportion preferring treatment A and, given θ, assume the participant responses are random and independent. Assume the following 10

responses are observed: $\mathbf{y} = (A, B, A, A, A, A, A, B, A, A)$. The likelihood function is then $L_{\mathbf{y}}(\theta) \propto \theta^8(1-\theta)^2$. Now assume two different stopping rules were used to collect $\mathbf{y}$: 1) stop sampling after 10 participant responses and 2) stop sampling once the absolute difference in the number of participant preferences for A and B is 6 (i.e., $|A - B| = 6$). Under either stopping rule, the posterior distribution for θ will depend on the data $\mathbf{y}$ only through $L_{\mathbf{y}}(\theta)$. Thus, if the same prior distribution for θ is used, Bayesian adaptive trials will yield the same posterior distribution regardless of the stopping rule used, remaining consistent with the likelihood principle.

Contrast this with a frequentist analysis of the same sequential experiment, employing the same stopping rules but analyzing the data using p-values. Recall the p-value is the probability under the null hypothesis of obtaining a result as extreme or more extreme than what is observed. Consider null and alternative hypotheses of the form $H_0 : \theta = 1/2$ and $H_1 : \theta > 1/2$ and let Y_i be an indicator function for participant i preferring A. Under stopping rule (1), the p-value is obtained as the tail probability of a Binomial($n = 10$, $\theta = 1/2$) distribution:

$$\text{p-value} = P\left(\sum_{i=1}^{10} Y_i \geq 8\right) = \left[\binom{10}{8} + \binom{10}{9} + \binom{10}{10}\right]\left(\frac{1}{2}\right)^{10} = \frac{56}{1024}.$$

Under stopping rule (2), the p-value is the probability under the null of 8 A's and 2 B's (i.e., what is observed), 7 A's and 1 B, or 6 A's in a row:

$$\text{p-value} = \left[\binom{6}{1}\binom{2}{1} + \binom{6}{2}\right]\left(\frac{1}{2}\right)^{10} + \binom{6}{1}\left(\frac{1}{2}\right)^{8} + \binom{6}{0}\left(\frac{1}{2}\right)^{6} = \frac{67}{1024}.$$

We see that different stopping rules lead to different inference under a frequentist p-value approach, a clear violation of the likelihood principle.

2.2.3 Adaptive stopping rules

The desire for designs with improved efficiency but which remained consistent with the likelihood principle motivated early work in adaptive stopping rules for Bayesian clinical trials.

Some of these designs were decision-theoretic in nature, and proposed determining the optimal (sequential) stopping rules as those which maximized the expected utility of the trial. Anscombe (1963) describes a decision-theoretic futility stopping rule for normally distributed data which explicitly considers the possibility of incorrectly treating future patients. Berry (1987) asks if, in addition to the possibility of incorrectly treating future patients, the possibility of incorrectly treating current patients in the trial should also be considered. Weighing these two objectives could be accomplished using a different decision-theoretic futility stopping rule. He notes, however, that determining optimal stopping rules for decision-theoretic designs involves challenging computations which require backward induction. Backward induction is an iterative procedure which infers an optimal sequence of decisions. It does so by determining the optimal decision at each time point, starting from the final decision and working backwards to the initial decision. An example of determining optimal stopping rules in this manner is provided in Berry and Ho (1988) for a decision-theoretic design for normally distributed data which includes a futility stopping rule and incorporates the profitability of a drug-development program into the utility function. For designs permitting many interim analyses (e.g., fully sequential designs), backward induction may be computationally infeasible. Thus, Carlin et al. (1998) proposed a computationally efficient forward sampling algorithm which permits identification of optimal stopping rules in decision-theoretic designs with a larger number of interim analyses.

For many designs which are not decision-theoretic in nature, the posterior distribution of parameters of interest or the posterior predictive distribution of future outcomes may be utilized to define adaptive stopping rules (Berry, 1987). Posterior predictive probabilities have been used for sequential decision making for many years (e.g., Thatcher (1964)). Herson (1979) proposed using them for futility stopping in a trial with dichotomous outcomes. He derived analytic solutions for the posterior predictive probabilities under a binomial likelihood and beta prior, and suggested stopping the trial once the probability of futility became too high. In a similar spirit, Berry (1985) considered an example of normally distributed

data where the posterior distribution of the mean parameter was used for early stopping. He showed that the expected sample size decreased as the frequency of interim analyses increased. In the following, I adapt the example he used to illustrate a stopping rule which is a functional of the posterior distribution.

Consider a trial where we are interested in assessing whether an unknown mean θ is different from 0. We thus have the following null and alternative hypotheses: $H_0 : \theta = 0$ and $H_1 : \theta \neq 0$. Assume the participant responses at an interim analysis performed at time t are $Y_i \overset{iid}{\sim} N(\theta, 1)$ for $i = 1, ..., n_t$ where $\overline{Y}_t = \frac{1}{n_t} \sum_{i=1}^{n_t} Y_i$. We assign the following prior, $\theta \sim N(0, \tau^2)$, where τ^2 denotes the prior variance. After observing $Y_1, ..., Y_{n_t}$, we have the posterior distribution for θ as:

$$p(\theta \mid \bar{y}_t) = N\left(\frac{n\bar{y}_t}{n + \tau^{-2}}, \frac{1}{n + \tau^{-2}}\right).$$

We must now decide whether to stop the trial or continue enrolling participants. Intuitively, the trial should stop if there is sufficient evidence to suggest that θ is different from 0. Formally, this occurs when $p(\theta < 0 \mid \bar{y}_t) > u_t$ or when $p(\theta > 0 \mid \bar{y}_t) > u_t$, where u_t is a probability threshold that reflects the strength of evidence required to stop the trial and may depend on t. Higher values of u_t correspond to requiring stronger evidence before stopping the trial. Thus, the trial stops if

$$p(\theta < 0 \mid \bar{y}_t) = \Phi\left(\frac{-n\bar{y}_t}{\sqrt{n + \tau^{-2}}}\right) > u_t \quad \text{or} \quad p(\theta > 0 \mid \bar{y}_t) = 1 - \Phi\left(\frac{-n\bar{y}_t}{\sqrt{n + \tau^{-2}}}\right) > u_t$$

and continues otherwise. Selection of an appropriate value for u_t depends on the objectives of the trial. It may be specified by subject matter experts, or potentially optimized to satisfy frequentist operating characteristics required by regulatory agencies. In this example, new participants would continue to be enrolled until either the stopping rule is satisfied at a future interim analysis, a maximum sample size is reached, or a summary of the posterior distribution attains some value (e.g., reaching a specified level of posterior precision).

More recent works have utilized posterior distribution-based early stopping to expand beyond the normal endpoint case. For example, Thall and Simon (1994) consider this approach for binary efficacy outcomes, Thall et al. (1995) consider it for multiple outcomes, such as efficacy and toxicity, and Thall et al. (2005) extend the approach to handle time-to-event endpoints. Indeed, adaptive stopping rules which are functionals of the posterior or posterior predictive distributions have become quite common and are now routinely recommended in practice (Berry et al., 2010).

2.3 Late phase Bayesian adaptive RCTs

The goal of late phase RCTs is to confirm whether or not an experimental treatment is efficacious. Many late phase Bayesian adaptive RCTs are sequentially randomized and include adaptive stopping rules (i.e., decision criteria), defined as posterior probability statements about an estimand. An estimand is a target of inference, and in the RCT setting, the estimand of interest is typically a marginal treatment effect. I denote the marginal estimand by $\gamma(\boldsymbol{\theta})$, which is a function of model parameters $\boldsymbol{\theta}$ from a correctly specified analysis model. Without loss of generality, assume the alternative hypothesis of the trial specifies the marginal treatment effect being greater than a clinically meaningful threshold, γ_0:

$$H_0 : \gamma(\boldsymbol{\theta}) \leq \gamma_0 \text{ vs } H_A : \gamma(\boldsymbol{\theta}) > \gamma_0.$$

I use $T(\mathbf{y}_{n_t})$ to denote the posterior probability that the alternative hypothesis is true given the data, $\mathbf{y}_{n_t}$, which is collected from the n_t participants who are enrolled in the trial at an interim or final analysis conducted at time t:

$$T(\mathbf{y}_{n_t}) = P(H_A|\mathbf{y}_{n_t}) = P(\gamma(\boldsymbol{\theta}) > \gamma_0|\mathbf{y}_{n_t}). \tag{2.3}$$

The decision criteria are defined based on this posterior probability to declare treatment superiority or futility at any interim or final analysis. As described in the last section, treatment superiority is declared if the posterior probability exceeds some upper probability

threshold, u_t, i.e., if $T(\mathbf{y}_{n_t}) > u_t$, or treatment futility is declared if the posterior probability falls below some lower probability threshold, l_t, i.e., if $T(\mathbf{y}_{n_t}) < l_t$. If a superiority or futility declaration is made at an interim analysis, the trial stops early.

The most common approach in the design of late phase Bayesian adaptive RCTs is a hybrid approach, where the upper and lower probability thresholds are set to $u_t = u$ and $l_t = l$ for all t, and are then optimized so that the trial design has desirable values of frequentist operating characteristics (e.g., Type 1 error rate and statistical power; Berry et al. (2010)). As $T(\mathbf{y}_{n_t})$ defined above is a functional of the posterior distribution of the marginal estimand $\gamma(\boldsymbol{\theta})$, improving its estimation may yield better trial operating characteristics.

2.3.1 Estimands and collapsibility

The marginal estimand, $\gamma(\boldsymbol{\theta})$, that is used to specify $T(\mathbf{y}_{n_t})$ in (2.3), can be defined as a contrast $f(\cdot)$ between population averaged quantities, $\mu(\boldsymbol{\theta}; A)$, which are functions of unknown parameters $\boldsymbol{\theta}$, and treatment assignments A. If we let $A = 1$ denote assignment to the treatment arm and $A = 0$ denote assignment to the control arm, then the marginal estimand can be defined as follows:

$$\gamma(\boldsymbol{\theta}) = f(\mu(\boldsymbol{\theta}; A = 1), \mu(\boldsymbol{\theta}; A = 0)). \tag{2.4}$$

As an example, $\gamma(\boldsymbol{\theta})$ might represent the difference in means of a continuous endpoint when comparing the treated to untreated populations. In that case, $f(\cdot)$ would be the difference operator, and $\mu(\boldsymbol{\theta}; A = 1)$ and $\mu(\boldsymbol{\theta}; A = 0)$ would represent the mean of that endpoint in those who are treated and untreated, respectively. We may also consider conditional parameters $\mu(\boldsymbol{\theta}; A, \mathbf{X})$ which contain additional covariates $\mathbf{X}$. Using the example above, $\mu(\boldsymbol{\theta}; A = 1, \mathbf{X} = \mathbf{x})$ would represent the mean of the continuous endpoint in those treated individuals who have covariate pattern $\mathbf{X} = \mathbf{x}$.

An estimand is said to be *collpasible* when it can be represented as a contrast between either

the marginal or conditional parameters μ for a fixed covariate pattern:

$$\begin{aligned}\gamma(\boldsymbol{\theta}) &= f(\mu(\boldsymbol{\theta}; A=1), \mu(\boldsymbol{\theta}; A=0)) \\ &= f(\mu(\boldsymbol{\theta}; A=1, \mathbf{X}=\mathbf{x}), \mu(\boldsymbol{\theta}; A=0, \mathbf{X}=\mathbf{x})).\end{aligned} \tag{2.5}$$

When collapsibility holds, samples from the posterior distribution of $\gamma(\boldsymbol{\theta})$ may be obtained using an analysis model which either includes or excludes $\mathbf{X}$. I refer to a model which includes $\mathbf{X}$ as an *adjusted* model, and one which excludes $\mathbf{X}$ as an *unadjusted* model. When the second equality in (2.5) does not hold, and the marginal treatment effect cannot be represented as a contrast of the conditional parameters, the estimand is said to be *non-collpasible* (e.g., Daniel et al. (2021)). Under non-collapsibility, samples from the posterior distribution of $\gamma(\boldsymbol{\theta})$ may still be obtained directly using an unadjusted model. However, posterior samples can no longer be obtained directly from adjusted models. This is the case for many estimands which are defined using non-linear endpoints, and which are commonly used within late phase RCTs (e.g., relative risk, odds ratio, hazard ratio). To obtain samples from the posterior of $\gamma(\boldsymbol{\theta})$, posterior samples from adjusted analyses must undergo a post-processing marginalization step, such as regression-based standardization (Kalton, 1968).

2.3.2 Covariate adjustment

Covariate adjustment can be defined as adjustment for baseline values of variables known to be associated with the outcome of interest (also referred to as "prognostic covariates"; United States Food and Drug Administration (2021)). In this subsection, I review the use of covariate adjustment for improving the estimation of treatment effects and the operating characteristics of RCTs. In the next subsection, I review a historical approach toward personalized medicine, where covariates are used to identify treatment-respondent subgroups (i.e., subgroup analysis). I focus on the case where adjustment is performed using linear models and generalized linear models only.

I begin by defining two correlated random variables, Y and Z, which are bivariate normally

distributed according to the following:

$$\begin{pmatrix} Y \\ Z \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_y \\ \mu_z \end{pmatrix}, \begin{pmatrix} \sigma_y^2 & \rho\sigma_y\sigma_z \\ \rho\sigma_y\sigma_z & \sigma_z^2 \end{pmatrix} \right). \tag{2.6}$$

By standard results of the multivariate normal distribution, we have the following conditional distribution for $Y \mid Z = z$:

$$Y \mid Z = z \sim N\left(\mu_y + \frac{\sigma_y}{\sigma_z}\rho(z - \mu_z), (1 - \rho^2)\sigma_y^2\right). \tag{2.7}$$

We can see that when we condition on the value of the correlated variable Z, there is a reduction in the variability of Y, the magnitude of which depends on the strength of correlation ρ. This result helps motivate the discussion below about covariate adjustment in RCTs with normal endpoints.

The use of covariate adjustment to improve the precision of estimates from a linear model dates back to at least the 1920's (Fisher, 1925). Cox (1957) describes how performing a linear regression on a single prognostic covariate improves the precision of an estimated treatment effect. Letting $Var(Y) = \sigma_y^2$ above represent the residual variance from a fitted linear regression which does not include Z, we see from (2.7) that including Z in the model reduces this residual variance to $Var(Y \mid Z = z) = (1 - \rho^2)\sigma_y^2$. Cox and McCullagh (1982) use this result to state that for randomized experiments under the assumption of no treatment covariate interactions, the primary purpose of covariate adjustment is to reduce the residual variance by a factor $(1-\rho^2)$ so that the precision of the treatment contrasts is increased. This increase in precision leads to increases in statistical power (Pocock et al., 2002). Senn (1989) considers a parallel-group RCT with a normal outcome and known variance and considers a single prognostic covariate, say Z. He focuses on testing the difference in means between the treatment groups and shows that, conditioned on a particular standardized imbalance between the treatment groups for Z, denoted by δ_z, an analysis adjusting for the covariate yields the nominal significance level. However, conditioned on the same δ_z, an unadjusted

analysis does not yield the nominal significance level. Rather, for an unadjusted analysis, the conditional size of the test depends on the sign of δ_z and on the strength of correlation between the outcome and Z. Pocock et al. (2002) elaborate on this example, showing in their Figure 2 that even when Z is perfectly balanced across treatments (i.e., $\delta_z = 0$), unadjusted tests are *conservative* (i.e., less than the nominal significance level) when large correlations between the outcome and Z exist. Thus, to increase precision and power, and to preserve the conditional size of the statistical tests employed, Senn (1989) argues it is prudent to adjust for covariates in trials with normal endpoints and that, as standardized imbalances are independent of sample size, adjustment should be made regardless of the sample size of the trial.

Covariate adjustment has also been considered for non-linear endpoints (see Hauck et al. (1998)). When looking at the precision of the estimators for covariate-adjusted treatment effects for these endpoints (e.g., the regression coefficient for a treatment indicator from generalized linear models which include additional covariates), there is actually a *decrease* in precision. Robinson and Jewell (1991) showed this for logistic regression models for binary endpoints, and Ford et al. (1995) showed this for Cox proportional hazards models when considering time-to-event endpoints. Paradoxically, this decrease in precision from covariate adjustment still results in *increased* power, as shown in Hernández et al. (2004) for binary endpoints and in Hernández et al. (2006) for time-to-event endpoints. This paradox is resolved, however, by noting that the standard error comparisons are being made between estimators of different estimands, which results from the non-collapsibility of treatment effects for non-linear endpoints (Daniel et al., 2021). Specifically, unadjusted analyses yield estimates of marginal estimands, and covariate-adjusted analyses yield estimates of conditional estimands. The power increase results from the value of the conditional estimand moving away from the null at a faster rate than the increase in the standard error of its estimator. However, as mentioned in the previous subsection, to obtain estimates of marginal estimands from adjusted analyses, marginalization (e.g., through standardization) is required.

Benkeser et al. (2021) show that when estimates obtained from covariate-adjusted analyses are marginalized using standardization, the covariate-adjusted analyses lead to both increased power and increased precision of the marginal treatment effect as compared to the unadjusted analyses.

Given the benefits described above, along with the fact that it is common in RCTs to collect a substantial amount of baseline covariate information (Friedman et al., 2015), both the FDA and European Medicines Agency now recommend covariate adjustment be performed in practice (United States Food and Drug Administration, 2021; European Medicines Agency, 2015). They emphasize that any covariates used for adjustment be pre-specificed and selected based on *a priori* subject matter expertise, a practice which Ciolino et al. (2019) showed has been increasing over time. Recent research has also demonstrated the benefits of covariate adjustment in more flexible frequentist designs. Van Lancker et al. (2022) considered its use in group sequential and information adaptive designs, showing increased statistical power results from adjusting for covariates at interim analyses in addition to the final analysis. Lee et al. (2022) investigated covariate adjustment in the context of adaptive multi-arm designs for marginal treatment effects defined using continuous endpoints, showing that in addition to increasing statistical power, covariate adjustment also increases the probability of selecting effective treatments for continuation in the later stages of the trial. Despite this recent interest, there has been no assessment of covariate adjustment in flexible Bayesian adaptive designs.

2.3.3 Personalized medicine

The FDA states the goal of personalized medicine is "to target the right treatments to the right patients at the right time" (United States Food and Drug Administration, 2018a). One historical approach toward personalized medicine has included performing subgroup analyses in late phase RCTs. This approach mainly consists of identifying significant treatment-covariate interactions in a fitted statistical model. In the context of parametric survival

models, Byar and Corle (1977) proposed identifying significant treatment-covariate interactions through the use of likelihood ratio tests. The presence of significant interactions leads to the possibility of a "treatment of choice", where optimal treatments may be identified based on patient characteristics. In their study, they defined a patient's optimal treatment as that which minimized the patient's estimated hazard function.

Simon (1982) considers subgroup analysis in the context of logistic regression and Cox proportional hazards models. He stresses the importance of pre-specifying a limited number of subgroups of interest to be tested via formal interaction tests, and cautions against performing hypothesis tests within subgroups since repeated testing increases the chance of false positive findings. He states that under a null hypothesis of no treatment effect, an upper bound on the probability of finding at least one subgroup with a significant treatment effect is $1-(1-\alpha)^G$, where α is the significance level of each test and G is the number of subgroup-specific hypothesis tests performed. This quantity represents an upper bound since the hypothesis tests are correlated, not independent. Despite this caution, Pocock et al. (2002) report a case where 24 subgroup analyses were performed in a single clinical trial. If each subgroup analysis was performed using a hypothesis test with significance level of $\alpha = 0.05$ under the null hypothesis of no treatment difference, then an upper bound on the probability of at least one subgroup displaying a significant treatment effect is $1-(1-0.05)^{24} = 0.71$. Thus, under this scenario, it is likely that a significant treatment effect would have been identified in at least one subgroup due to chance alone. Formal interaction tests attempt to safeguard against this.

Gail and Simon (1985) describe different types of treatment-covariate interactions that may exist in clinical trials. They define a qualitative (or crossover) interaction to be present when the directions of the true treatment differences vary across subgroups, though they state interactions of this type are thought to be rare in practice. If present, however, they have important implications for ethical reasons. A qualitative interaction also includes the

case of there being a response in some subgroups and no response in others (Russek-Cohen and Simon, 1997). A non-crossover interaction is then defined to be present when the directions of the treatment differences are the same for all subgroups, but the magnitudes of the differences differ across subgroups. Finally, a quantitative interaction is defined to be the existence of any heterogeneity of treatment differences across subgroups. Gail and Simon (1985) then continue by proposing a likelihood ratio test to identify qualitative interactions and also describe how to test for quantitative interactions.

It is important to note that clinical trials are often powered to detect marginal treatment effects only, and so formal interaction tests are generally underpowered. While sometimes this lack of power is used as an argument against formal interaction tests, Pocock et al. (2002) argue this is actually a benefit of their usage, as interaction tests recognize the limits inherent in the collected data and guard against exaggerated claims of subgroup differences. To further guard against exaggerated claims of subgroup differences, Bayesian approaches for subgroup analysis have been proposed. When no *a priori* differences are expected among the patient subgroups with respect to treatment, Dixon and Simon (1991) show how an assumption of exchangeable treatment-covariate interaction terms shrinks their posterior distributions toward zero, thereby reducing the possibility of exaggerated subgroup differences. Additionally, Simon (2002) shows how a Bayesian approach for subgroup analysis can incorporate the *a priori* belief that qualitative interactions are rare in practice. The goals of personalized medicine may also be pursued within the context of early phase dose-finding trials, which I review next.

2.4 Early phase Bayesian adaptive dose-finding trials

The goal of early phase clinical trials is to assess the safety and efficacy of first-in-human doses of experimental treatments and to identify doses which should be further investigated.

Early phase Bayesian adaptive designs are commonly employed for these trials. These so-called "dose-finding" designs can generally be divided into three major families according to increasing model complexity: algorithmic, model-assisted, and model-based (Berry et al., 2010; Yuan et al., 2019). Algorithmic designs require no statistical modeling and perform dose-finding using a pre-defined algorithm. Perhaps the best known of these designs is the "$3+3$" design, which uses cohorts of three patients at each iteration of the algorithm to find a maximum tolerated dose (MTD), defined as the largest dose which produces a toxicity response closest to a pre-defined threshold (Storer, 1989). Model-assisted designs are more complex and use a statistical model to pre-compute decision boundaries which help define the dose-finding procedure. These designs offer better statistical properties than algorithmic designs while maintaining their simplicity. One of the better known model-assisted designs is the "Bayesian optimal interval" design (Yuan et al., 2016a). At each iteration of dose-finding, this design calculates a statistic using the accumulating trial data and compares it to pre-defined decision boundaries to determine whether to maintain, escalate, or de-escalate the current dose. As compared to the $3+3$ design, it has been shown to select the MTD more often and assign more patients to it as well. While algorithmic and model-assisted designs are easy to implement, they are less flexible and can be less efficient than model-based designs, and I do not consider them further in this thesis.

Among dose-finding methods, model-based designs offer the greatest flexibility and have superior statistical properties, but they are also the most challenging to implement. They require the specification of a full probability model for the dose-response surfaces of interest, which permits inclusion of prior information if available. Incorporating prior information can be especially beneficial in early phase designs since they typically have very small sample sizes. As mentioned previously, prior information may come in the form of clinical experience, previous animal or clinical studies, or from pharmacokinetic/pharmacodynamic models. For these designs, the dose-finding procedure is defined with respect to the fitted model, with designs being differentiated by the dose-response model they propose and by the procedure

they use to select the next dose for evaluation. A large amount of historical research in early phase model-based designs focused on the development of cytotoxic agents in oncology, where interest was in identifying the MTD using toxicity responses only. Under the assumption of monotonically increasing dose-efficacy and dose-toxicity surfaces, the MTD is also the most efficacious dose. Many of these designs focused exclusively on binary toxicity endpoints for monotherapy, where only a single cytotoxic agent is of interest. For example, O'Quigley et al. (1990) proposed the Continual Reassessment Method which utilizes accumulating patient responses to sequentially assign doses to new patients estimated to be closest to the MTD. Babb et al. (1998) proposed the sequential Escalation with Overdose Control design, which seeks to identify the MTD while also limiting the number of patients being assigned to overly toxic doses.

2.4.1 Dose-finding for continuous endpoints

Interest in using continuous endpoints for early phase dose-finding trials has grown over time. Le Tourneau et al. (2009) describe how there is increasing use of molecularly targeted agents as treatments, e.g., immunotherapies, which may provide a therapeutic response at dosages which are lower than the MTD. In trials investigating these compounds, toxicity responses are replaced by measures of biological activity, e.g., plasma drug concentration or measures of target inhibition in tissues of interest, which are often measured on the continuous scale (Korn, 2004; Le Tourneau et al., 2009). Furthermore, some have proposed using continuous toxicity scores (e.g., Chen et al. (2010); Lee et al. (2012)) in lieu of binary toxicity responses. Le-Rademacher et al. (2020) propose a continuous Adverse Event (AE) burden score which is defined using both frequency and severity of AEs experienced by a patient throughout the trial. They argue it offers a richer and more nuanced understanding of patients' total toxicity burden in the trial than would using a binary toxicity response. Despite this increasing interest, there have been fewer designs proposed for continuous responses than for binary responses. The Continual Reassessment Method and the Escalation with Overdose Control

designs have been extended to this setting through the use of quasi-likelihood approaches (Yuan et al., 2007; Chen et al., 2012). Other designs have been proposed which assume monotononically non-decreasing continuous dose-response surfaces (e.g., Eichhorn and Zacks (1973); Wang and Ivanova (2015)).

2.4.2 Optimal biological dose combinations

With the increasing use of moleculary targeted agents where the MTD may not be the most efficacious dose, focus is placed on identifying an optimal biological dose (OBD) by using designs which incorporate both efficacy and toxicity (e.g., Thall and Cook (2004); Yuan et al. (2016b)). These designs have been extended to the combination therapy setting, where two or more agents are combined and determining an OBD *combination* is of interest. Wages and Conaway (2014) proposed a design which seeks to identify the most efficacious dual-agent dose combination which is below a specified toxicity level, under the assumption of monotone dose-response surfaces when the levels of one agent is fixed. While the assumption of monotonicity may be reasonable in some cases, it does not hold in general (Li et al., 2017). To identify OBD combinations under possible non-monotonic dose-response surfaces, flexible designs have been proposed (e.g., Mozgunov and Jaki (2019b)). Houede et al. (2010) propose a flexible design to optimize a utility function over ordinal efficacy and toxicity responses. They use a generalized Aranda-Ordaz model (Aranda-Ordaz, 1981) which is extremely flexible. However, the model is highly parameterized and may introduce estimation challenges given the small sample sizes in early phase trials. Flexible designs for continuous dose-response surfaces have also been proposed for combination therapies (e.g., Mozgunov and Jaki (2019a)).

2.4.3 Personalized dose-finding for combination therapies

Many dose-finding designs assume that the optimal dose is the same for every patient in the population. This includes all designs previously described above. I refer to this approach as *standard dose-finding*. Contrast this with *personalized dose-finding*, whose goal is to identify

patient-specific optimal doses. This latter approach toward dose-finding better aligns with the aim of personalized medicine, which was previously described as being "to target the right treatments to the right patients at the right time" (United States Food and Drug Administration, 2018a). Despite regulatory support for this aspiration, the field of personalized dose-finding trials remains underdeveloped since it is difficult to extend many parametric dose-finding methods to this setting, as potentially many dose-covariate interaction terms must be estimated. These trials are even more challenging in the combination therapy setting, where the additional dosing agents amplify the estimation challenges. One of the few examples of personalized dose-finding for combination therapies is Mozgunov et al. (2022), who proposed a dual-agent personalized dose-finding design which was motivated by an opiate detoxification trial. However, only a single dosing agent was optimized in this design, as the level of the other agent was pre-specified externally by clinicians.

Personalized dose-finding for combination therapies can be viewed as an optimization problem. Let $\mathbf{Z} = \{Z_p\}_{p=1}^{P}$ be a set of discrete covariates used for tailoring, where K strata are defined as the Cartesian product of their levels, and let $\mathbf{d} = (d_1, ..., d_J) \in \mathbb{D} \subset \mathbb{R}^J$ be a dose combination combining J dosing agents. Let $f(\mathbf{d}, \mathbf{z})$ be a continuous efficacy function of interest and $g(\mathbf{d}, \mathbf{z})$ be a continuous toxicity function, where both are transformed such that smaller values denote being more desirable. Then personalized dose-finding seeks to find an OBD combination for each of the K strata. Specifically, the goal is to find, for each stratum k, the $\mathbf{d}$ which minimizes $f(\mathbf{d}, \mathbf{z}_k)$, subject to being below a tolerable toxicity threshold $g_k^\dagger$, yielding the following *constrained* optimization problem:

$$\underset{\mathbf{d} \in \mathbb{D}}{\operatorname{argmin}} f(\mathbf{d}, \mathbf{z}_k) \quad \text{subject to} \quad g(\mathbf{d}, \mathbf{z}_k) \leq g_k^\dagger \text{ for } \quad k = 1, ..., K. \tag{2.8}$$

Stratum-specific toxicity thresholds $g_k^\dagger$ allow tolerable toxicity levels to depend on specific covariate profiles. This may be useful in settings where the clinical team permits greater levels of toxicity for those patients with more severe variants of a disease, for example. Note that the optimization problem above could be defined with respect to a utility function

which considers both $f(\mathbf{d}, \mathbf{z})$ and $g(\mathbf{d}, \mathbf{z})$, though this approach is not considered further in this thesis. When no drug-related adverse events are expected and a minimal toxicity setting can be assumed, toxicity information does not need to be incorporated into the dose-finding. Under this setting, $g(\mathbf{d}, \mathbf{z}_k)$ is dropped from the optimization problem in (2.8) creating an *unconstrained* optimization problem over the efficacy function $f(\mathbf{d}, \mathbf{z}_k)$ only:

$$\underset{\mathbf{d} \in \mathbb{D}}{\operatorname{argmin}} f(\mathbf{d}, \mathbf{z}_k) \quad \text{for} \quad k = 1, ..., K. \tag{2.9}$$

Any strategy proposed to solve the optimization problems in (2.8) and (2.9) must accommodate the small sample sizes and high levels of observation noise which are characteristic of early phase designs.

2.5 Bayesian optimization

Bayesian optimization is a derivative-free method which finds the global optima of expensive-to-evaluate objective functions (Garnett, 2023). It utilizes stochastic surrogate models, commonly GPs, to estimate the objective function, and then defines a sequential search policy through use of a so-called acquisition function. It is sample efficient, seeking to identify the optima while at the same time minimizing the number of objective function evaluations required to do so. For these reasons, Bayesian optimization is a particularly attractive method for dose-finding problems, where objective function evaluations require a patient to be given an experimental therapy and potentially carry both a high financial and ethical cost. Other global optimization methods exist (e.g., genetic algorithms and simulated annealing), but require many objective function evaluations or gradient information and so are not appropriate for dose-finding (Bull, 2011; Tracey and Wolpert, 2018).

Bayesian optimization has long been used for global optimization problems in the engineering, machine learning, and computer experiments literature (e.g., Kushner (1964); Zhilinskas (1975); Jones et al. (1998); Gramacy (2020)). Recently within the biostatistics literature, some have proposed using it to efficiently find optimal dynamic treatment regimes (e.g.,

Rodriguez Duque et al. (2022); Freeman et al. (2022)). Furthermore, it has received some attention in the early phase Bayesian adaptive dose-finding literature. For example, Takahashi and Suzuki (2021b) proposed using Bayesian optimization to find the MTD in early phase clinical trials for monotherapies, focusing exclusively on binary toxicity responses. Using a GP to estimate the logit transformed probability of toxicity at a particular dose, they defined the dose-finding strategy using the Expected Improvement acquisition function (Jones et al., 1998), to be further described below. The same authors also proposed using Bayesian optimization to identify monotherapy OBDs by considering both binary efficacy and binary toxicity responses (Takahashi and Suzuki, 2021a). They optimized a utility function which was defined using the efficacy and toxicity response probabilities, and employed the Lower Confidence Bound acquisition function (Srinivas et al., 2010) to perform the dose-finding.

2.5.1 Gaussian process regression

A GP is a stochastic process, any finite subsample of which is distributed as multivariate normal. A GP can be viewed as a prior distribution over the function space, yielding an analytically tractable posterior distribution over functional forms which are explained by the observed data (Williams and Rasmussen, 2006). To model the efficacy function, a GP prior is used:

$$f(\mathbf{d}, \mathbf{z}) \sim GP\left(m(\mathbf{d}, \mathbf{z}), \nu\mathcal{K}_{\boldsymbol{\eta}}((\mathbf{d}, \mathbf{z}), (\mathbf{d}', \mathbf{z}'))\right) \tag{2.10}$$

where $m(\mathbf{d}, \mathbf{z})$ is the mean function, and $\mathcal{K}_{\boldsymbol{\eta}}((\mathbf{d}, \mathbf{z}), (\mathbf{d}', \mathbf{z}'))$ is a correlation function (kernel) parameterized by $\boldsymbol{\eta}$ and multiplied by scale parameter ν (Binois and Gramacy, 2021). The scale parameter determines the variability of the efficacy function throughout the domain. It is common to specify a constant mean function, $m(\mathbf{d}, \mathbf{z}) = \beta$, since the correlation function is responsible for much of the method's flexibility (Gramacy, 2020).

Using the GP prior as specified above induces a multivariate normal distribution on the observations, $\mathbf{y} \sim N(\beta\mathbf{1}_n, \nu\mathbf{K}_{\boldsymbol{\eta}})$, where $\mathbf{K}_{\boldsymbol{\eta}}(i, j) = \mathcal{K}_{\boldsymbol{\eta}}((\mathbf{d}_i, \mathbf{z}_i), (\mathbf{d}_j, \mathbf{z}_j)) + \tau^2\mathbb{1}_{i=j}$ and τ^2

represents observation noise. While taking a fully Bayesian approach is possible by specifying prior distributions for the kernel hyperparameters $\boldsymbol{\theta} = \{\nu, \tau^2, \boldsymbol{\eta}\}$ and sampling from their joint posterior distribution, it is common to proceed in an empirical Bayes fashion since it is less computationally demanding, an important consideration within the context of sequential optimization (Gramacy, 2020). Thus, after data $\mathcal{D} = \{(\mathbf{d}_i, \mathbf{z}_i, y_i)\}_{i=1}^n$ are observed, the kernel hyperparameters are estimated via maximum likelihood and replaced by their point estimates. This yields the posterior distribution of the therapeutic function at a new dose combination $\widetilde{\mathbf{d}}$ in stratum k, denoted by $\widetilde{\mathbf{d}}_k = (\widetilde{\mathbf{d}}, \widetilde{\mathbf{z}}_k)$, as $p(f \mid \mathcal{D}, \widetilde{\mathbf{d}}_k) = N(\mu(\widetilde{\mathbf{d}}_k), \sigma^2(\widetilde{\mathbf{d}}_k))$ (Binois and Gramacy, 2021), such that

$$\begin{aligned}
\mu(\widetilde{\mathbf{d}}_k) &= \widehat{\beta} + \mathbf{k}_{\boldsymbol{\eta}}(\widetilde{\mathbf{d}}_k)^T \mathbf{K}_{\boldsymbol{\eta}}^{-1}(\mathbf{y} - \widehat{\beta}\mathbf{1}) \\
\sigma^2(\widetilde{\mathbf{d}}_k) &= \nu\mathcal{K}_{\boldsymbol{\eta}}(\widetilde{\mathbf{d}}_k, \widetilde{\mathbf{d}}_k) - \nu\mathbf{k}_{\boldsymbol{\eta}}(\widetilde{\mathbf{d}}_k)^T \mathbf{K}_{\boldsymbol{\eta}}^{-1}\mathbf{k}_{\boldsymbol{\eta}}(\widetilde{\mathbf{d}}_k) + \frac{(1 - \mathbf{k}_{\boldsymbol{\eta}}(\widetilde{\mathbf{d}}_k)^T \mathbf{K}_{\boldsymbol{\eta}}^{-1}\mathbf{1})^2}{\mathbf{1}^T \mathbf{K}_{\boldsymbol{\eta}}^{-1}\mathbf{1}} \\
\widehat{\beta} &= \frac{\mathbf{1}^T \mathbf{K}_{\boldsymbol{\eta}}^{-1}\mathbf{y}}{\mathbf{1}^T \mathbf{K}_{\boldsymbol{\eta}}^{-1}\mathbf{1}}
\end{aligned} \tag{2.11}$$

where $\mathbf{K}_{\boldsymbol{\eta}}$ is $n \times n$, $\mathbf{k}_{\boldsymbol{\eta}}(\widetilde{\mathbf{d}}_k) = [\mathcal{K}_{\boldsymbol{\eta}}((\mathbf{d}_1, \mathbf{z}_1), \widetilde{\mathbf{d}}_k), ..., \mathcal{K}_{\boldsymbol{\eta}}((\mathbf{d}_n, \mathbf{z}_n), \widetilde{\mathbf{d}}_k)]^T$ is $n \times 1$, and $\widehat{\beta}$ is a plug-in estimate for the mean. This posterior distribution is used to define the acquisition functions which control the sequential search strategy of the dose-finding methods.

2.5.2 Acquisition functions

An acquisition function, denoted by $\alpha(\widetilde{\mathbf{d}}_k \mid \mathcal{D})$, defines the search policy used within a Bayesian optimization framework. After the initial data are collected for c_k dose combinations and an initial GP model fitted, the next dose combination for evaluation within stratum k is selected as the maximizer of $\alpha(\widetilde{\mathbf{d}}_k \mid \mathcal{D})$:

$$\mathbf{d}_k^{(c_k+1)} = \underset{\widetilde{\mathbf{d}} \in \mathbb{D}}{\operatorname{argmax}}\, \alpha(\widetilde{\mathbf{d}}_k \mid \mathcal{D}). \tag{2.12}$$

Many commonly used acquisition functions have a basis in Bayesian decision theory and can be viewed as performing one-step lookahead over an implicitly defined utility function for the data. That is, if the next dose combination is the last to be collected, then many commonly

used acquisition functions can be viewed as selecting that dose combination which maximizes the expected marginal gain in utility of the dataset (see Table 7.1 in Garnett (2023)). Having a basis in Bayesian decision theory provides many theoretical performance guarantees, such as convergence to the global optima under certain conditions (Bull, 2011).

One of the major challenges in dose-finding is adequately exploring the dose-response surface while at the same time identifying desirable regions where the OBD combination may lie. If local optima exist, search policies may become stuck and never converge to the global optima. The Expected Improvement, denoted by $\alpha_{EI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D})$, is perhaps the most commonly employed acquisition function for unconstrained optimization problems like (2.9) and balances this exploration-exploitation trade-off. Intuitively, α_{EI} quantifies how much improvement the dose combination $\widetilde{\mathbf{d}}_k$ is expected to yield over the current best observation of $f(\mathbf{d}, \mathbf{z}_k)$, denoted by f_k^*, and is defined as the following expectation:

$$\alpha_{EI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D}) = \mathbb{E}\left[\max(0, f_k^* - f(\widetilde{\mathbf{d}}_k)) \mid \mathcal{D}, \widetilde{\mathbf{d}}_k\right]. \tag{2.13}$$

Under a GP posterior, Jones et al. (1998) showed this expectation to be analytically tractable, which is an important consideration given the computational demands of the overall optimization scheme:

$$\alpha_{EI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D}) = (f_k^* - \mu(\widetilde{\mathbf{d}}_k))\Phi\left(\frac{f_k^* - \mu(\widetilde{\mathbf{d}}_k)}{\sigma(\widetilde{\mathbf{d}}_k)}\right) + \sigma(\widetilde{\mathbf{d}}_k)\phi\left(\frac{f_k^* - \mu(\widetilde{\mathbf{d}}_k)}{\sigma(\widetilde{\mathbf{d}}_k)}\right). \tag{2.14}$$

In this equation, $\Phi(\cdot)$ and $\phi(\cdot)$ denote the standard normal cumulative distribution function and probability density function, respectively, and $\mu(\widetilde{\mathbf{d}}_k)$ and $\sigma(\widetilde{\mathbf{d}}_k)$ are given by (2.11). The second term gives large values to dose combinations for which $f(\mathbf{d}, \mathbf{z}_k)$ is imprecisely estimated (i.e., exploration) and the first term gives large values to dose combinations for which $f(\mathbf{d}, \mathbf{z}_k)$ is desirable (i.e., exploitation).

To consider dose-finding under higher-grade toxicities as in (2.8), α_{EI} can be extended to the constrained Expected Improvement, denoted by $\alpha_{cEI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D})$, where toxicity information in

the form of $g(\mathbf{d}, \mathbf{z}_k)$ is formally incorporated into the expectation:

$$\alpha_{cEI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D}) = \mathbb{E}\left[\max(0, f_k^* - f(\widetilde{\mathbf{d}}_k))\mathbb{1}\{g(\widetilde{\mathbf{d}}_k) \leq g_k^\dagger\} \mid \mathcal{D}, \widetilde{\mathbf{d}}_k\right]. \tag{2.15}$$

In many contexts it may be reasonable to assume that the efficacy and toxicity functions are conditionally independent given dose. If each response surface is modeled using independent GPs, then α_{cEI} is also available analytically (Gardner et al., 2014):

$$\alpha_{cEI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D}) = \alpha_{EI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D})\Phi\left(\frac{g_k^\dagger - \mu_g(\widetilde{\mathbf{d}}_k)}{\sigma_g(\widetilde{\mathbf{d}}_k)}\right) \tag{2.16}$$

where $\mu_g(\widetilde{\mathbf{d}}_k)$ and $\sigma_g(\widetilde{\mathbf{d}}_k)$ represent the posterior mean and standard deviation of the toxicity function at $\widetilde{\mathbf{d}}_k$, taking the same functional forms as in (2.11). The exploration-exploitation trade-off defined by α_{EI} is now weighted by the posterior probability of satisfying the toxicity constraint $g_k^\dagger$, giving higher weight to dose combinations which are more likely to be safe.

2.6 Summary

In this literature review, I introduced the core themes which are present throughout this thesis as well as reviewed the historical and related work which motivate it. I began by introducing clinical trials, and then narrowed my focus to a subgroup of these which form the main area of interest of this thesis, Bayesian adaptive clinical trials. I described how the decision criteria which govern the adaptations in these trials are often functionals of the posterior distribution, and how Bayesian trials which employ early stopping rules satisfy the likelihood principle. I next reviewed how covariate adjustment has been used to improve trial operating characteristics and to permit a personalized approach toward medicine in late phase Bayesian adaptive RCTs. I then shifted focus to early phase designs, where I described historical approaches toward dose-finding and the challenges of evaluating combination therapies in a personalized medicine setting. In a final section, I discussed Bayesian optimization within the context of early phase dose-finding.

Chapter 3

Covariate adjustment in Bayesian adaptive randomized controlled trials

Preamble to Manuscript 1. Covariate adjustment has been shown to increase statistical power for fixed size, late phase RCTs (e.g., Benkeser et al. (2021)). Recently, its impact in more flexible frequentist designs has been assessed. Van Lancker et al. (2022) showed increases in statistical power for group sequential and information adaptive designs which utilize covariate adjustment at interim analyses in addition to the final analysis. Lee et al. (2022) noted similar power increases in adaptive multi-arm designs for marginal treatment effects defined using continuous endpoints, noting that covariate adjustment also increases the probability of selecting effective treatments. Despite this recent interest, there has been no investigation of covariate adjustment in flexible Bayesian adaptive designs, including the impact of combining it with prior information on the covariate effects.

In this manuscript, covariate adjustment is investigated for Bayesian adaptive designs which permit early stopping for treatment superiority. The impact of adjustment is assessed for several data generating mechanisms, which include continuous, binary, and time-to-event outcomes, for trials of different maximum sample sizes. We consider several forms of ad-

justment model misspecification and investigate the impact of using prior information of varying strengths. Our interest is in marginal treatment effects, and we provide a tutorial for obtaining marginal posterior samples from adjusted analyses. The original contributions of this chapter are i) investigating the impact of covariate adjustment in Bayesian adaptive clinical trials, ii) assessing its impact under a wider range of scenarios than previously reported for flexible designs, and iii) combining covariate adjustment with prior information on the covariate effects. The substantive contribution of this work is providing further theoretical support for the benefits of covariate adjustment and describing how adjustment can be performed in practice for Bayesian adaptive clinical trials.

The corresponding manuscript was published in *Statistical Methods in Medical Research* (Willard et al., 2024).

Covariate Adjustment in Bayesian Adaptive Randomized Controlled Trials

James Willard[1], Shirin Golchi[1], and Erica E.M. Moodie[1]

[1]*Department of Epidemiology, Biostatistics, and Occupational Health, McGill University*

This thesis contains the version of the corresponding paper published in *Statistical Methods in Medical Research* (Willard et al., 2024).

Abstract

In conventional randomized controlled trials, adjustment for baseline values of covariates known to be at least moderately associated with the outcome increases the power of the trial. Recent work has shown particular benefit for more flexible frequentist designs, such as information adaptive and adaptive multi-arm designs. However, covariate adjustment has not been characterized within the more flexible Bayesian adaptive designs, despite their growing popularity. We focus on a subclass of these which allow for early stopping at an interim analysis given evidence of treatment superiority. We consider both collapsible and non-collapsible estimands, and show how to obtain posterior samples of marginal estimands from adjusted analyses. We describe several estimands for three common outcome types. We perform a simulation study to assess the impact of covariate adjustment using a variety of adjustment models in several different scenarios. This is followed by a real world application of the compared approaches to a COVID-19 trial with a binary endpoint. For all scenarios, it is shown that covariate adjustment increases power and the probability of stopping the trials early, and decreases the expected sample sizes as compared to unadjusted analyses.

3.1 Introduction

In conventional, fixed size randomized controlled trials (RCTs), adjustment for baseline values of covariates known to be at least moderately associated with the outcome has been shown to increase the power of the trial (Hernández et al., 2004; Benkeser et al., 2021; Hernández et al., 2006; Harrell and Slaughter, 2021). This is because covariate adjustment improves the precision of the estimated treatment effect and accounts for the outcome heterogeneity within each treatment arm that is explained by the adjustment variables (Lewis, 1999; Senn, 2013; Benkeser et al., 2021). Adjustment also corrects for any chance imbalance of important baseline variables which may exist post-randomization (Kahan et al., 2014). Therefore, covariate adjustment in the primary analysis of clinical trials is now recommended by both the US Food and Drug Administration (FDA) and European Medicines Agency (EMA) (United States Food and Drug Administration, 2021; European Medicines Agency, 2015). Additionally, systematic reviews have suggested its use in practice has grown over time (Austin et al., 2010; Ciolino et al., 2019). Recently, these power increases have been demonstrated in more flexible frequentist designs, such as information adaptive and adaptive multi-arm designs (Van Lancker et al., 2022; Lee et al., 2022). However, the simulation scenarios investigated in each of these designs contained at least one of the following: only continuous outcomes with no treatment-covariate interactions where the marginal and conditional estimands are the same; only a single sample size; a data generating process containing only a small number of variables; or only a small number of covariate adjustment models. A more comprehensive investigation is needed to better understand the benefits of covariate adjustment under a broader array of flexible design scenarios.

While the impact of covariate adjustment has been demonstrated in the flexible frequentist designs mentioned above, it has not been characterized within flexible Bayesian adaptive designs, where early stopping at an interim analysis is permitted given evidence of treatment superiority or futility. In these designs one may learn about a treatment effect while

potentially requiring fewer participants than fixed designs. Additionally, Bayesian trials allow for seamless incorporation of prior information for model parameters including covariate effects. The impact of combining prior information with covariate adjustment has not been previously investigated. With the growing interest in Bayesian adaptive designs, a characterization of covariate adjustment for several commonly used sample sizes and outcome types would be highly valuable for researchers.

In this work, we consider the impact of covariate adjustment in Bayesian adaptive designs which allow for early stopping for superiority and provide a step-by-step tutorial for post adjustment marginalization. We explore several data generating processes for continuous, binary, and time-to-event outcomes, and consider adjustment models which include several forms of misspecification while incorporating varying levels of prior information for the covariate effects. The covariate adjustment described herein is performed using generalized linear models (GLMs). However, the methods, results, and recommendations discussed below are not specific to these models, and are expected to generalize to other parametric and nonparametric models.

The manuscript is organized in the following manner. We first introduce and describe Bayesian adaptive designs which allow for early stopping, as well as the targets of inference (i.e., estimands), which are marginal treatment effects for each endpoint. We then describe the specific collapsible and non-collapsible estimands used in this manuscript. For the non-collapsible estimands, we describe their estimation and marginalization through a Bayesian framework. This is followed by a simulation study which shows the impact of covariate adjustment on design operating characteristics including power, the probability of stopping the trial early, and expected sample size, for multiple sample sizes. We then show a real-world application of covariate adjustment within a COVID-19 RCT and end with a discussion.

3.2 Bayesian adaptive designs with early stopping

Bayesian adaptive designs allow for predetermined changes to the trial design at interim analyses based on evidence provided by the accumulating data (Berry et al., 2010; Giovagnoli, 2021). These designs include sequentially randomized trials which allow for early stopping for superiority or futility at interim analyses. Interim analyses are performed to determine whether to stop the trial early and declare treatment superiority or futility, or to continue the trial. The decision to stop the trial early at an interim analysis is controlled by a predefined decision rule. Decision criteria may be defined with respect to different functionals of the posterior distribution of the parameter of interest or estimand. Posterior or posterior predictive probability statements about the estimand are commonly used statistics. In the RCT setting, estimands are typically defined as marginal treatment effects in a specified population of interest. We adopt this convention throughout, but delay further discussion of marginal estimands and their posterior estimation until the next section.

In the Bayesian adaptive designs described in this work, interim or final decisions are defined with respect to posterior probability statements about a marginal treatment effect, $\gamma(\boldsymbol{\theta})$, which is a function of model parameters $\boldsymbol{\theta}$. The alternative hypothesis of the trial is formulated as this marginal treatment effect being greater than a clinically meaningful threshold, γ_0:

$$H_0 : \gamma(\boldsymbol{\theta}) \leq \gamma_0 \text{ vs } H_A : \gamma(\boldsymbol{\theta}) > \gamma_0.$$

A Bayesian test statistic can be defined as the posterior probability of this alternative hypothesis given the data, $\mathcal{D}_{n_t} = \{Y_{n_t}, A_{n_t}, \mathbf{X}_{n_t \times p}\}$, which may include any observed outcomes (Y_{n_t}), treatment assignments (A_{n_t}), and p additional covariates ($\mathbf{X}_{n_t \times p}$), for the n_t participants who are enrolled in the trial at an interim or final analysis conducted at time t:

$$T(\mathcal{D}_{n_t}) = P(H_A|\mathcal{D}_{n_t}) = P(\gamma(\boldsymbol{\theta}) > \gamma_0|\mathcal{D}_{n_t}). \tag{3.1}$$

This statistic is then used to define a decision rule, which declares treatment superiority at

any interim or final analysis if the statistic exceeds some upper probability threshold, u, i.e., if

$$T(\mathcal{D}_{n_t}) > u.$$

If a superiority declaration is made at an interim analysis, the trial stops early. The common approach in the design of Bayesian adaptive trials is the "hybrid" approach, where the upper probability threshold, u, is optimized so that the trial design has desirable values of frequentist operating characteristics (Berry et al., 2010). For example, power ($\mathcal{P}$) and the Type 1 error rate (T1E) are defined as follows:

$$\mathcal{P} = P(T(\mathcal{D}_{n_t}) > u \mid \gamma(\boldsymbol{\theta}) = \gamma^* > \gamma_0)$$

$$\text{T1E} = P(T(\mathcal{D}_{n_t}) > u \mid \gamma(\boldsymbol{\theta}) = \gamma_0).$$

Since the sampling distribution of a Bayesian posterior probability is generally unknown, calibration of the design to meet frequentist operating characteristics requires simulation studies. Note that models that adjust for covariates under most settings result in analytically intractable posterior distributions. Therefore, the evaluation of $T(D_{n_t})$ within every trial simulation requires posterior sampling or approximation techniques. In this paper we use Markov chain Monte Carlo (MCMC) to sample from the posterior distribution when not available in closed form.

3.3 Estimands and Bayesian estimation

In what follows, let A_i be defined as a binary treatment assignment for the i^{th} participant, where $A_i = 1$ represents being randomized to the treatment group and $A_i = 0$ to the control group. Let $\widetilde{\mathbf{X}}_{n_t \times d}$ represent a matrix of $j = 1, ..., d$ covariates measured at baseline for $i = 1, ..., n_t$ participants in the study at an interim or final analysis conducted at time t. Let $\tilde{\mathbf{x}}_i = (x_1, x_2, ..., x_d)_i$ represent the row of this matrix corresponding to the full covariate pattern of the i^{th} participant, and let Y_i be an arbitrary outcome of interest for the i^{th} participant. For notational simplicity, subscripts are dropped for the remainder of this article

except when strictly necessary. Treatments are assigned through simple randomization and are thus independent of all covariates measured at baseline (i.e., $A \perp\!\!\!\perp \tilde{\mathbf{x}}$), but we allow for the possibility of chance imbalance of covariates between treatment groups.

We use the term "unadjusted analysis" to refer to a model which includes only the treatment assignment indicator (A), and the term "adjusted analysis" for a model which includes $p \leq d$ additional covariates, $\mathbf{X} \subseteq \widetilde{\mathbf{X}}$, with $\mathbf{x}_i = (x_1, ..., x_p)_i$ representing the row of $\mathbf{X}$ corresponding to the i^{th} participant's covariate pattern used for adjustment. For adjusted analyses, many different covariate sets may be adjusted for in addition to the binary treatment indicator. This includes any covariates $\mathbf{Z} \subseteq \mathbf{X}$ where a treatment-covariate interaction exists. Let ϕ be the regression coefficient for the treatment assignment indicator, β_0 be the intercept, $\boldsymbol{\beta} = \{\beta_1, ..., \beta_p\}$ be the vector of covariate main effects, and $\boldsymbol{\omega} = \{\omega_1, ..., \omega_m\}$ be the vector of treatment-covariate interaction effects for those covariates used in the adjustment model within a GLM setting. Additionally, let $\boldsymbol{\zeta}$ be a set of nuisance parameters not of direct interest but which are required for model specification (e.g., baseline hazard parameters in a time-to-event setting). Then let $\boldsymbol{\theta} = \{\beta_0, \phi, \boldsymbol{\beta}, \boldsymbol{\omega}, \boldsymbol{\zeta}\}$ be the set of model parameters with prior $p(\boldsymbol{\theta})$, and define $\eta(\boldsymbol{\theta}; A_i, \mathbf{X}_i, \mathbf{Z}_i)$ to be the expected outcome for participant i on the linear scale. Let $p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta})$ represent each participant's contribution to the likelihood function and $g(\cdot)$ be the link function. Assuming independence among participants, we have the joint posterior distribution of the model parameters, $\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t})$, under arbitrary adjustment model specification, being proportional to the likelihood times the prior:

$$\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t}) \propto \prod_{i=1}^{n_t} p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta}) p(\boldsymbol{\theta}).$$

In this manuscript, the decision rule used at an interim or final analysis is defined with respect to a posterior probability statement about a marginal estimand $\gamma(\boldsymbol{\theta})$, presented in Equation 3.1, which is the average treatment effect in a population of interest. We note this marginal estimand can be defined as a contrast of population average quantities, such as a difference of means or a ratio of population-level event or survival probabilities. Let $\mu(\boldsymbol{\theta}; A)$

be the population level parameter for treatment group A used as input for contrast $f(\cdot)$. Then $\gamma(\boldsymbol{\theta})$ can be represented by:

$$\gamma(\boldsymbol{\theta}) = f(\mu(\boldsymbol{\theta}; A = 1), \mu(\boldsymbol{\theta}; A = 0)).$$

Unadjusted analyses yield posterior samples from the marginal parameter $\mu(\boldsymbol{\theta}; A)$ directly:

$$\mu(\boldsymbol{\theta}; A) = g^{-1}(\eta(\boldsymbol{\theta}; A)).$$

Adjusted analyses, however, yield posterior samples from the conditional parameter $\mu(\boldsymbol{\theta}; A, \mathbf{X})$ for fixed covariate pattern $\mathbf{X}$:

$$\mu(\boldsymbol{\theta}; A, \mathbf{X}) = g^{-1}(\eta(\boldsymbol{\theta}; A, \mathbf{X}, \mathbf{Z})).$$

When a treatment effect is *collapsible*, it can be represented as a contrast between either the marginal or conditional parameters for a fixed covariate pattern $\mathbf{X}$:

$$\begin{aligned}\gamma(\boldsymbol{\theta}) &= f(\mu(\boldsymbol{\theta}; A = 1), \mu(\boldsymbol{\theta}; A = 0)) \\ &= f(\mu(\boldsymbol{\theta}; A = 1, \mathbf{X}), \mu(\boldsymbol{\theta}; A = 0, \mathbf{X})).\end{aligned}$$

Thus, under collapsibility, samples from the posterior distribution of the marginal estimand can be obtained from either an unadjusted or adjusted analysis. When a treatment effect is *non-collapsible*, the marginal treatment effect cannot be represented as a contrast between conditional parameters for a fixed covariate pattern $\mathbf{X}$ (Daniel et al., 2021). This is commonly the case for treatment effects modeled by GLMs or in the presence of treatment-covariate interactions. As an example, consider a hypothetical RCT with a binary endpoint which follows the following logistic regression model with binary treatment assignment A and binary covariate X, where $P(X = 1) = P(A = 1) = 0.5$ and where $\phi = \log(5)$, $\beta = \log(10)$, and $\boldsymbol{\theta} = \{\phi, \beta\}$:

$$\text{logit}(P(Y = 1 \mid A, X)) = \phi A + \beta X.$$

Define $\mu(\boldsymbol{\theta}; A, X)$ to be the treatment specific conditional risk. The conditional odds ratio for

those who are treated versus untreated can be represented as a contrast of these conditional risk parameters, and is 5 regardless of the value of X. To find the marginal odds ratio, the treatment specific conditional risks must be averaged with respect to the distribution of X before calculating the odds ratio, effectively collapsing over X in a stratified two-by-two table. Doing so gives treatment specific marginal risks $\mu(\boldsymbol{\theta}; A)$ which are used to obtain a marginal odds ratio of 4.1. This value is smaller than that for the conditional odds ratio and shows the marginal odds ratio cannot be represented as a contrast of the treatment specific conditional risks. Thus, the odds ratio is non-collapsible (see Section A.7 of Appendix A for more details). Under non-collapsibility, samples from the posterior distribution of the marginal estimand can still be obtained directly from unadjusted analyses, but not from adjusted analyses. To obtain samples from the posterior distribution of the marginal estimand using an adjusted analysis, the posterior samples of $\mu(\boldsymbol{\theta}; A, \mathbf{X})$ must be marginalized with respect to the distribution of $\mathbf{X}$, yielding samples of $\mu(\boldsymbol{\theta}; A)$ which are then used in the contrast:

$$\mu(\boldsymbol{\theta}; A) = \int_{\mathbf{X}} \mu(\boldsymbol{\theta}; A, \mathbf{X}) p(\mathbf{X}) d\mathbf{X}.$$

This marginalization is commonly called standardization or Bayesian G-computation for point treatments (Kalton, 1968; Freeman Jr and Holford, 1980; Lane and Nelder, 1982; Saarela et al., 2015a; Remiro-Azócar et al., 2022; Keil et al., 2018; Daniel et al., 2021). The integral over $p(\mathbf{X})$ is approximated through summation where, for each of $s = 1, ..., S$ Monte Carlo samples $\boldsymbol{\theta}_s$ from $\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t})$, the following calculations is performed:

$$\mu(\boldsymbol{\theta}_s; A) \approx \sum_{i=1}^{n_t} w_{i,s} \mu(\boldsymbol{\theta}_s; A, \mathbf{x}_i). \tag{3.2}$$

The $\mathbf{x}_i$ are the covariate patterns used for adjustment and which are contained in the joint empirical distribution of the collected sample data. The weights $\mathbf{w}_s = (w_{1,s}, ..., w_{n_t,s})$ are sampled as $\mathbf{w}_s \sim \text{Dirichlet}(\mathbf{1}_{n_t})$, corresponding to the Bayesian bootstrap, where $\mathbf{1}_{n_t}$ is the n_t-dimensional vector of 1's (Rubin, 1981; Oganisian and Roy, 2021; Linero and Antonelli,

2023). This can then be used to obtain a single sample from the posterior of the marginal treatment effect:

$$\gamma(\boldsymbol{\theta}_s) = f(\mu(\boldsymbol{\theta}_s; A = 1), \mu(\boldsymbol{\theta}_s; A = 0)).$$

Note that when any X_j for $j = 1, ..., p$ is the propensity score being jointly modeled with the outcome of interest, a different Bayesian bootstrap procedure should be utilized (Stephens et al., 2023). For the remainder of the manuscript, it is assumed $\mathbf{X}$ does not contain a propensity score estimated in this manner.

The procedure above enables S samples from the posterior distribution of the marginal estimand $\gamma(\boldsymbol{\theta})$ to be obtained using an arbitrary adjustment model. This allows for a direct performance comparison between adjustment models within Bayesian adaptive designs. Below we describe this procedure within the context of specific models that are most commonly used for different outcome types in clinical trials.

3.3.1 Collapsible treatment effects

Difference in means: no treatment-covariate interactions

Consider the difference in means of a continuous endpoint under the assumption of no treatment-covariate interactions (i.e., $\mathbf{Z} = \emptyset$; homogeneity of the treatment effect),

$$\gamma(\boldsymbol{\theta}) := \mu(\boldsymbol{\theta}; A = 1) - \mu(\boldsymbol{\theta}; A = 0)$$

where $\mu(\boldsymbol{\theta}; A = a) = E[Y \mid A = a; \boldsymbol{\theta}]$. This marginal estimand represents the difference in expected outcomes between those assigned to treatment versus those assigned to control. Estimation proceeds assuming independent outcomes and the following model:

$$\begin{aligned}
&p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta}) = \text{Normal}(\mu(\boldsymbol{\theta}; A_i, \mathbf{X}_i), \sigma^2) \\
&\mu(\boldsymbol{\theta}; A_i, \mathbf{X}_i) = \beta_0 + \phi A_i + \mathbf{X}_i\boldsymbol{\beta} \\
&\boldsymbol{\theta} = \{\beta_0, \phi, \boldsymbol{\beta}, \sigma^2\}.
\end{aligned}$$

Since this treatment effect is collapsible, posterior samples of the marginal estimand can be obtained from either an unadjusted or adjusted analysis, and samples from the posterior distribution of the treatment indicator coefficient ϕ are commonly used.

3.3.2 Non-collapsible treatment effects

Difference in means: treatment-covariate interactions

Consider again a continuous outcome, but under the assumption of at least one treatment-covariate interaction (i.e., $\mathbf{Z} \neq \emptyset$; treatment effect heterogeneity). The difference in means is non-collapsible. Estimation proceeds assuming independent outcomes and the following model:

$$p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta}) = \text{Normal}(\mu(\boldsymbol{\theta}; A_i, \mathbf{X}_i), \sigma^2)$$

$$\mu(\boldsymbol{\theta}; A_i, \mathbf{X}_i) = \beta_0 + \phi A_i + \mathbf{X_i}\boldsymbol{\beta} + (A_i \cdot \mathbf{Z_i})\boldsymbol{\omega}$$

$$\boldsymbol{\theta} = \{\beta_0, \phi, \boldsymbol{\beta}, \boldsymbol{\omega}, \sigma^2\}.$$

Since this estimand is non-collapsible, posterior samples of the conditional $\mu(\boldsymbol{\theta}; A, \mathbf{X})$ must be marginalized using (3.2) before forming the contrast to obtain a posterior sample of the marginal difference in means. An outline of this procedure is provided in Section A.1 in Appendix A.

Relative risk and odds ratio

Consider a dichotomous outcome which is modelled as a Bernoulli random variable. Examples of commonly used marginal estimands include the relative risk

$$\gamma(\boldsymbol{\theta}) := \mu(\boldsymbol{\theta}; A = 1)/\mu(\boldsymbol{\theta}; A = 0)$$

and the odds ratio

$$\gamma(\boldsymbol{\theta}) := \frac{\mu(\boldsymbol{\theta}; A = 1)/(1 - \mu(\boldsymbol{\theta}; A = 1))}{\mu(\boldsymbol{\theta}; A = 0)/(1 - \mu(\boldsymbol{\theta}; A = 0))}$$

where $\mu(\boldsymbol{\theta}; A = a) = E[Y \mid A = a; \boldsymbol{\theta}]$. The relative risk represents the ratio comparing the risk of an event for those assigned to treatment versus those assigned to control. The odds

ratio represents the ratio comparing the odds of an event for those assigned to treatment versus those assigned to control. Estimation proceeds assuming independent outcomes and the following model:

$$
\begin{aligned}
&p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta}) = \text{Bernoulli}(\mu(\boldsymbol{\theta}; A_i, \mathbf{X}_i)) \\
&\mu(\boldsymbol{\theta}; A_i, \mathbf{X}_i) = \text{logit}^{-1}(\beta_0 + \phi A_i + \mathbf{X}_i\boldsymbol{\beta} + (A_i \cdot \mathbf{Z}_i)\boldsymbol{\omega}) \\
&\boldsymbol{\theta} = \{\beta_0, \phi, \boldsymbol{\beta}, \boldsymbol{\omega}\}.
\end{aligned}
$$

To obtain posterior samples of the marginal relative risk or odds ratio, posterior samples of the conditional $\mu(\boldsymbol{\theta}; A, \mathbf{X})$ must be marginalized using (3.2) before forming the contrast, an outline of which is provided in Section A.1 in Appendix A.

Hazard ratio

Let T denote the time to an event of interest. Let $h(t \mid A)$ represent the hazard, the instantaneous event rate at time t, for those assigned to treatment A:

$$
h(t \mid A) = \lim_{\Delta t \to 0} \frac{P(t \leq T < t + \Delta t \mid T > t, A)}{\Delta t}.
$$

Under the assumption of no competing risks, there is a one-to-one relationship between the hazard and survival probability at time t:

$$
S(t \mid A) = \exp\left(-\int_{v=0}^{t} h(v \mid A) dv\right).
$$

Further assuming proportional hazards, an estimand of interest is the marginal hazard ratio

$$
\gamma(\boldsymbol{\theta}) := \log\{\mu(\boldsymbol{\theta}; A = 1)\} / \log\{\mu(\boldsymbol{\theta}; A = 0)\}
$$

where $\mu(\boldsymbol{\theta}; A = a) = S(t \mid A = a; \boldsymbol{\theta})$. This estimand represents the ratio comparing the hazard of those assigned to treatment versus those assigned to control. We note that the estimation framework described below is general and may be utilized to target other estimands of interest (e.g., the risk difference or risk ratio).

Under an RCT framework which allows for right censoring only, the time from a participant's initial enrollment to an event of interest may occur after the trial has ended. Let T_i be the i^{th} participant's observed event time or right censoring time. Let δ_i be the i^{th} participant's observation indicator, where $\delta_i = 1$ means the event time is observed before the end of the trial and where $\delta_i = 0$ means the event time is right censored. Let $Y_i = \{T_i, \delta_i\}$ be the observed data. On the hazard scale, the hazard of the event at time t for the i^{th} participant can be modeled as below, where $h_0(t)$ is the baseline hazard function:

$$h_i(t \mid A_i, \mathbf{X}_i) = h_0(t) \exp(\eta_i)$$
$$\eta_i = \phi A_i + \mathbf{X}_i \boldsymbol{\beta} + (A_i \cdot \mathbf{Z_i}) \boldsymbol{\omega}.$$

This yields the corresponding survival probability:

$$\begin{aligned} S_i(t \mid A_i, \mathbf{X}_i) &= \exp\left(- \int_{v=0}^{t} h_i(v \mid A_i, \mathbf{X}_i) dv \right) \\ &= \exp\left(- \int_{v=0}^{t} h_0(v) \exp(\eta_i) dv \right). \end{aligned}$$

The baseline hazard function may be flexibly modeled, with one possible choice being through M-splines (Brilleman et al., 2020). Let $M(t; \boldsymbol{\psi}, \boldsymbol{k}, \delta)$ be an M-spline function:

$$M(t; \boldsymbol{\psi}, \boldsymbol{k}, \delta) = \sum_{l=1}^{L} \psi_l M_l(t; \boldsymbol{k}, \delta).$$

Here $\boldsymbol{\psi}$ is the vector of coefficients for the L M-spline basis terms, with degree δ and knot locations $\boldsymbol{k}$. Integrating this M-spline function yields the following I-spline function, which is evaluated using the same coefficients, degree and knot locations:

$$I(t; \boldsymbol{\psi}, \boldsymbol{k}, \delta) = \sum_{l=1}^{L} \psi_l I_l(t; \boldsymbol{k}, \delta).$$

Both M-spline and I-spline functions can be evaluated analytically (Wang and Yan, 2021). By flexibly modeling the baseline hazard with M-splines, the hazard and the survival probability

become, respectively:

$$
\begin{aligned}
h_i(t \mid A_i, \mathbf{X}_i) &= M(t; \boldsymbol{\psi}, \boldsymbol{k}, \delta) \exp(\eta_i) \\
S_i(t \mid A_i, \mathbf{X}_i) &= \exp\left(- \int_{v=0}^{t} M(v; \boldsymbol{\psi}, \boldsymbol{k}, \delta) \exp(\eta_i) dv \right) \\
&= \exp\left(-I(t; \boldsymbol{\psi}, \boldsymbol{k}, \delta) \exp(\eta_i) \right).
\end{aligned}
$$

Estimation then proceeds by assuming independent outcomes and the following model:

$$
\begin{aligned}
p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta}) &= S_i(T_i \mid A_i, \mathbf{X}_i)^{1-\delta_i} h_i(T_i \mid A_i, \mathbf{X}_i)^{\delta_i} \\
\boldsymbol{\theta} &= \{\boldsymbol{\psi}, \phi, \boldsymbol{\beta}, \boldsymbol{\omega}\}.
\end{aligned}
$$

Posterior samples of the marginal hazard ratio are obtained by first marginalizing samples of $\mu(\boldsymbol{\theta}; A, \mathbf{X}) = S(t \mid A = a, \mathbf{X}; \boldsymbol{\theta})$ using (3.2) and then forming the contrast (Daniels et al., 2023; Stitelman et al., 2011; Remiro-Azócar et al., 2022). This procedure is outlined in Section A.1 in Appendix A.

3.4 Simulation study

In this section we perform simulations for the design and models described in the previous sections. We consider a design with a superiority stopping rule where superiority is declared at any interim or final analysis performed at time t if $T(\mathcal{D}_{n_t}) > 0.99$. The same value of $u = 0.99$ is selected for all maximum sample sizes to control the overall Type 1 error rate of the unadjusted model (i.e., Type 1 error rate below 0.05). The unadjusted model is selected as a conservative choice for trial planning purposes, since the true strength of any covariate effects and adjustment benefit may not be known in practice at the trial planning stage (Benkeser et al., 2021). Note that our interest is in comparing the performance of different adjustment models, not optimizing u for each maximum sample size and model, so a single conservative value of u above is chosen for all simulations. Our marginal treatment effects of interest are the difference in means of a Normal endpoint under the assumption of no treatment-covariate interactions, the relative risk under a binary endpoint, and the hazard

ratio under a time-to-event endpoint. Data generating processes with five covariates and a treatment assignment indicator are used for multiple sample sizes with each endpoint. We consider adjustment models which include several forms of misspecification and which incorporate varying levels of prior information for the covariate effects. To obtain marginal estimates from the adjustment models, the procedures described in the previous section are utilized. We follow a setup similar to that reported in a previous study which investigated covariate adjustment for endpoints commonly used in COVID-19 RCTs (Benkeser et al., 2021): for each maximum sample size with each endpoint, three treatment effect values are chosen. The first is the null treatment effect, and the second and third are those where the unadjusted model achieves roughly 50% and 80% power. This excludes the scenarios of a maximum sample size of 100 under the binary and time-to-event endpoints, whose second and third treatment effect sizes are chosen as those where the unadjusted model achieves roughly 30% or 40% power. This ensures all simulations maintain realistic values for the marginal treatment effects, and that the impact of covariate adjustment is compared at a value of the treatment effect for which trials are commonly powered (i.e., 80%). For each maximum sample size with each endpoint, the impact of covariate adjustment is quantified through the values of the following design operating characteristics: power, Type 1 error rate, expected sample size, probability of stopping early, bias and root mean squared error (RMSE).

3.4.1 Data generating mechanisms

For each combination of endpoint and maximum sample size $\{100, 200, 500, 1000\}$, the data generating mechanisms for the treatment assignment and covariate distributions measured at baseline are shown below, where joint independence between all variables is assumed. Letting η represent the linear predictor used in the data generating mechanisms, the set $\{\boldsymbol{\beta}, \phi\}$ represents the conditional covariate and treatment effects on the linear predictor

scale:

$$\eta = \beta_0 + \phi A + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_3^2 + \beta_5 X_5$$

$$\{A, X_1, X_2, X_6\} \sim \text{Bernoulli}(0.5)$$

$$\{X_3, X_5, X_7, X_8\} \sim \text{Normal}(0, 1)$$

$$\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5)$$

$$\gamma_{\text{max ss}} = \mathcal{S}(\eta)$$

$$\boldsymbol{\Theta} = \{(\gamma_{100} \times \boldsymbol{\beta}), (\gamma_{200} \times \boldsymbol{\beta}), (\gamma_{500} \times \boldsymbol{\beta}), (\gamma_{1000} \times \boldsymbol{\beta})\}.$$

Note that the variables $\{X_6, X_7, X_8\}$ are noise and are not predictive of, or correlated with, any other variables in the data generating mechanism. For the binary endpoint, β_0 is optimized to generate datasets which exhibit the correct marginal control event risk of $p_{ctr} = 0.3$ (Section A.2 in Appendix A). For the continuous and time-to-event endpoints, $\beta_0 = 0$. For the time-to-event endpoint, an exponential baseline hazard with rate $\lambda = 0.02$ is used. For the non-collapsible treatment effects, the true values of the marginal estimand $\gamma = \gamma(\boldsymbol{\theta}) = f(\mu(\boldsymbol{\theta}; A = 1), \mu(\boldsymbol{\theta}; A = 0))$ do not equal the conditional treatment effects $\phi = f(\mu(\boldsymbol{\theta}; A = 1, \mathbf{X}), \mu(\boldsymbol{\theta}; A = 0, \mathbf{X}))$ for fixed $\mathbf{X}$. Thus, the reported values of the marginal estimands are obtained through simulation (denoted by $\mathcal{S}(\cdot)$; Section A.2 in Appendix A), and the values of γ and ϕ are reported together. For the continuous endpoint, $\boldsymbol{\beta} = (0, 0.5, -0.25, 0.5, -0.05, 0.25)$. For the binary endpoint, $\boldsymbol{\beta} = (-1.26, 1, -0.5, 1, -0.1, 0.5)$. For the time-to-event endpoint, $\boldsymbol{\beta} = (0, 1, -0.5, 1, -0.1, 0.5)$. For each maximum sample size (max ss) within each outcome type, 1,000 treatment-covariate datasets are generated. These are used to generate 1,000 different outcome vectors for each value of the marginal treatment effect within the corresponding maximum sample size. The specific parameter values used for all simulations are included in Table 3.1.

3.4.2 Adjustment models

Six adjustment models are considered for all endpoints:

1. correct: $\beta_0 + \phi A + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_3^2 + \beta_5 X_5$

Table 3.1: Simulation study parameter settings for the marginal estimand (γ) and conditional treatment effect (ϕ) for each endpoint and maximum sample size (max ss). The marginal estimands for the continuous, binary, and time-to-event are, respectively, the difference in means under the assumption of no treatment-covariate interactions, the relative risk, and the hazard ratio.

Continuous

max ss	γ	max ss	γ
100	0	500	0
100	-0.52	500	-0.22
100	-0.73	500	-0.32
200	0	1000	0
200	-0.36	1000	-0.16
200	-0.52	1000	-0.22

Binary

max ss	γ	ϕ	max ss	γ	ϕ
100	1	0	500	1	0
100	0.53	-0.99	500	0.72	-0.56
100	0.46	-1.21	500	0.60	-0.82
200	1	0	1000	1	0
200	0.59	-0.86	1000	0.80	-0.39
200	0.41	-1.36	1000	0.72	-0.54

Time-to-event

max ss	γ	ϕ	max ss	γ	ϕ
100	1	0	500	1	0
100	0.65	-0.68	500	0.78	-0.39
100	0.60	-0.79	500	0.71	-0.54
200	1	0	1000	1	0
200	0.69	-0.59	1000	0.85	-0.27
200	0.57	-0.86	1000	0.78	-0.39

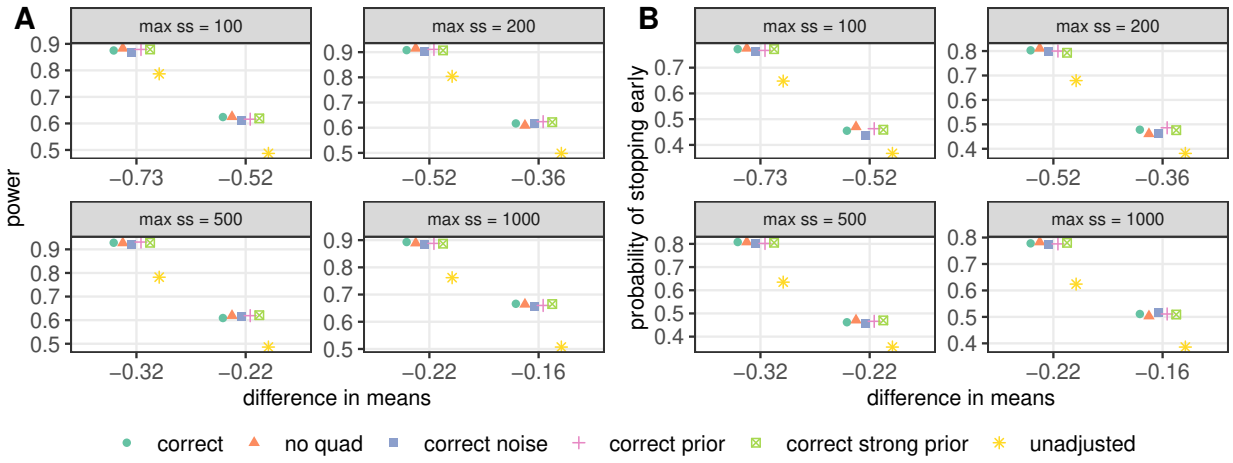

Figure 3.1: Continuous outcome. A) Power and B) probability of stopping early. Panels correspond to various maximum sample sizes (max ss). Points are jittered horizontally.

2. no quad: $\beta_0 + \phi A + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_5 X_5$

3. correct noise: $\beta_0 + \phi A + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_3^2 + \beta_5 X_5 + \beta_6 X_6 + \beta_7 X_7 + \beta_8 X_8$

4. correct prior: $\beta_0 + \phi A + \beta_1^\dagger X_1 + \beta_2^\dagger X_2 + \beta_3^\dagger X_3 + \beta_4^\dagger X_3^2 + \beta_5^\dagger X_5$

5. correct strong prior: $\beta_0 + \phi A + \beta_1^{\dagger\dagger} X_1 + \beta_2^{\dagger\dagger} X_2 + \beta_3^{\dagger\dagger} X_3 + \beta_4^{\dagger\dagger} X_3^2 + \beta_5^{\dagger\dagger} X_5$

6. unadjusted: $\beta_0 + \phi A$.

The *correct* model corresponds to an adjustment model which matches the data generating mechanism. The *no quad* model drops the quadratic component of X_3 from the *correct* model. The *correct noise* model adds noise variables $\{X_6, X_7, X_8\}$ to the *correct* model. These three models include priors for all parameters which are weakly informative only. The *correct prior* model is the same as the *correct* model, but includes priors for the covariate effects centered at the values used in the data generating mechanism. Similarly, the *correct strong prior* model both centers and re-scales these priors to be more informative. Note that the prior for the treatment indicator coefficient remains weakly informative in these models. Finally, the *unadjusted* model includes only the binary treatment indicator and uses weakly informative priors.

Table 3.2: Continuous outcome. Type 1 error rate (T1E), bias under the null (Bias*), and expected sample size at three different values of the marginal difference in means (γ).

	Maximum sample size = 100					Maximum sample size = 200				
			Expected sample size					Expected sample size		
Adjustment model	T1E	Bias*	$\gamma = 0$	$\gamma = -0.52$	$\gamma = -0.73$	T1E	Bias*	$\gamma = 0$	$\gamma = -0.36$	$\gamma = -0.52$
correct	0.034	-0.013	98.3	78.6	62.5	0.039	-0.011	196.8	154.7	116.0
no quad	0.037	-0.015	98.2	78.0	62.5	0.032	-0.009	197.5	156.2	115.4
correct noise	0.031	-0.012	98.6	80.3	64.8	0.036	-0.012	196.9	156.2	118.7
correct prior	0.032	-0.012	98.5	78.3	62.9	0.043	-0.012	196.4	154.5	116.3
correct strong prior	0.033	-0.012	98.3	78.4	62.5	0.041	-0.012	196.6	155.0	116.8
unadjusted	0.028	-0.007	98.5	82.6	68.9	0.039	-0.017	196.5	163.1	130.7
	Maximum sample size = 500					Maximum sample size = 1000				
			Expected sample size					Expected sample size		
Adjustment model	T1E	Bias*	$\gamma = 0$	$\gamma = -0.22$	$\gamma = -0.32$	T1E	Bias*	$\gamma = 0$	$\gamma = -0.16$	$\gamma = -0.22$
correct	0.036	-0.008	491.1	389.1	290.4	0.031	-0.003	987.0	753.0	598.8
no quad	0.031	-0.007	493.0	388.8	291.4	0.028	-0.003	987.8	755.2	597.5
correct noise	0.032	-0.006	492.8	391.1	293.0	0.030	-0.003	987.8	753.2	599.0
correct prior	0.032	-0.007	492.8	388.5	293.8	0.031	-0.003	987.2	753.5	598.0
correct strong prior	0.032	-0.007	492.8	388.8	291.0	0.029	-0.003	986.8	754.8	599.0
unadjusted	0.032	-0.006	492.4	415.6	343.1	0.028	-0.002	989.0	817.8	690.2

All simulations are performed using R (version 4.2.1). All modeling is performed using the GLM and survival functionality of the rstanarm package (version 2.21.2), a front-end to the STAN probabilistic programming language (R Core Team, 2022; Goodrich et al., 2020; Brilleman et al., 2020). For all coefficients other than $\{\beta_j^\dagger, \beta_j^{\dagger\dagger}\}$, the package's default weakly informative priors are used (Gabry and Goodrich, 2022). These priors induce moderate regularization and help improve computational stability. The prior for the intercept is placed after all covariates have been internally centered by rstanarm. Under a Normal likelihood for the continuous endpoint, this equates to the following priors, where $\beta_{0,c}$ represents the intercept's prior after covariate centering has been performed:

$$\begin{aligned} \beta_{0,c} &\sim \text{Normal}(\bar{y}, 2.5 s_y) & \beta_j^\dagger &\sim \text{Normal}(\beta_j, 2.5(s_y/s_x)) \\ \phi &\sim \text{Normal}(0, 2.5(s_y/s_x)) & \beta_j^{\dagger\dagger} &\sim \text{Normal}(\beta_j, s_y/s_x) \\ \beta_j &\sim \text{Normal}(0, 2.5(s_y/s_x)) & \sigma &\sim \text{Exponential}(1/s_y). \end{aligned}$$

Under the binary and time-to-event endpoints, the above priors for $\{\phi, \boldsymbol{\beta}\}$ are used with $\bar{y} = 0$ and $s_y = 1$. For the time-to-event endpoint, the coefficients of the M-spline basis ($\boldsymbol{\psi}$)

are constrained to a simplex to ensure identifiability of both the basis and linear predictor intercepts. Thus, the basis coefficients receive rstanarm's default Dirichlet prior (Brilleman et al., 2020). All models specify three Markov chains, each with 2,000 posterior samples. Half of the samples within each chain are used during the warm-up period, so 3,000 posterior samples in total are available for inference. Given the scale of the simulations performed, visual diagnostics assessing convergence of the Markov chains are not performed. Rather, for all simulations, values of STAN's implementation of the Gelman-Rubin $\hat{R}$ statistic are assessed to ensure Markov chain convergence (Gelman and Rubin, 1992).

The following null and alternative hypotheses are specified for the continuous, binary and time-to-event (TTE) endpoints, where $\gamma(\boldsymbol{\theta})$ is the marginal difference in means, marginal relative risk, and marginal hazard ratio, respectively:

$$\text{Continuous: } H_0 : \gamma(\boldsymbol{\theta}) \geq 0 \text{ vs } H_A : \gamma(\boldsymbol{\theta}) < 0$$

$$\text{Binary and TTE: } H_0 : \gamma(\boldsymbol{\theta}) \geq 1 \text{ vs } H_A : \gamma(\boldsymbol{\theta}) < 1.$$

For the continuous and binary endpoints, all outcomes are assumed to be observed immediately upon participant enrollment. For the time-to-event endpoint, it is assumed all outcomes are observed strictly after enrollment. For the continuous endpoint, interim analyses are performed after every 25, 50, 125, and 250 participants are enrolled for maximum samples sizes of 100, 200, 500, and 1,000, respectively. For the binary and time-to-event endpoints, interim analyses are event driven. For the binary endpoint, interim analyses are performed after at least 10, 20, 50, and 100 new events occur for maximum sample sizes 100, 200, 500, and 1,000, respectively. For the time-to-event endpoint, interim analyses are performed after at least 20, 40, 100, and 200 new events occur for maximum sample sizes 100, 200, 500, and 1,000, respectively. These numbers are chosen for each endpoint to ensure that on average the total number of analyses performed under the null treatment effect is less than four, which helps control the Type 1 error rate. They are also large enough to ensure there is a moderate chance of stopping at an early interim analysis under the non-null treatment

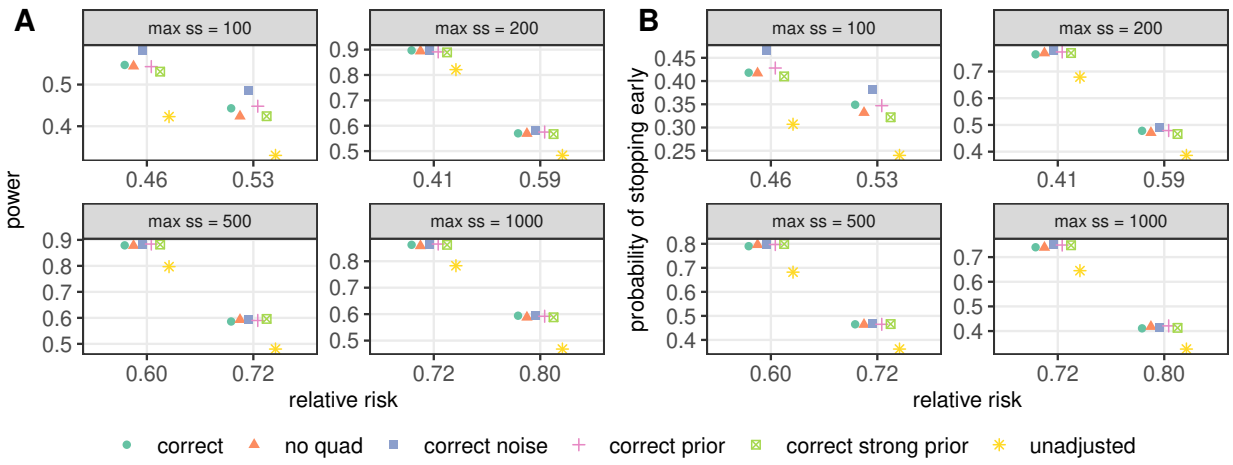

Figure 3.2: Binary outcome. A) Power and B) probability of stopping early. Panels correspond to various maximum sample sizes (max ss). Points are jittered horizontally.

effects. For the continuous and binary endpoints, interim analyses are performed until the trial is stopped early for superiority or until the maximum sample size is reached, at which time the final analysis is performed. For the time-to-event endpoint, interim analyses are performed until the trial is stopped early for superiority or until 50 time units from the start of the trial is reached, at which time the final analysis is performed. For this endpoint, participant enrollment is permitted until 25 time units. This ensures that participants enrolled at later time points are under follow-up long enough to have a moderately high probability of experiencing the event before the end of the trial. It also ensures that the maximum number of participants will be enrolled if there is not clear evidence of superiority at an early interim analysis. No loss to follow-up is assumed and administrative censoring of those still at risk is performed at 50 time units from the start of the trial.

3.4.3 Simulation study results

Within each outcome type and maximum sample size, the following design operating characteristics are investigated: power, Type 1 error rate, probability of stopping early, expected sample size, posterior median bias, and RMSE. Since the sampling distribution of the test statistic $T(\mathcal{D}_{n_t})$ is unknown, power and the Type 1 error rate are estimated via Monte Carlo using the 1,000 datasets. While this number is lower than that required by the FDA for

Table 3.3: Binary outcome. Type 1 error rate (T1E), bias under the null (Bias*), and expected sample size at three different values of the marginal relative risk (γ).

	Maximum sample size = 100					Maximum sample size = 200				
			Expected sample size					Expected sample size		
Adjustment model	T1E	Bias*	$\gamma = 1$	$\gamma = 0.53$	$\gamma = 0.46$	T1E	Bias*	$\gamma = 1$	$\gamma = 0.59$	$\gamma = 0.41$
correct	0.063	0.031	97.0	86.3	83.6	0.036	0.021	196.8	162.5	138.3
no quad	0.053	0.035	97.4	86.9	83.7	0.031	0.023	197.6	164.0	138.6
correct noise	0.079	0.031	96.2	84.2	81.0	0.044	0.021	196.0	161.3	137.4
correct prior	0.060	0.033	97.2	86.3	82.9	0.036	0.021	196.5	162.9	138.0
correct strong prior	0.052	0.035	97.7	88.1	84.5	0.036	0.022	196.9	164.6	138.2
unadjusted	0.034	0.058	98.6	90.7	88.3	0.031	0.025	198.0	171.4	147.2
	Maximum sample size = 500					Maximum sample size = 1000				
			Expected sample size					Expected sample size		
Adjustment model	T1E	Bias*	$\gamma = 1$	$\gamma = 0.72$	$\gamma = 0.60$	T1E	Bias*	$\gamma = 1$	$\gamma = 0.80$	$\gamma = 0.72$
correct	0.028	0.010	493.7	404.9	334.9	0.022	0.008	992.0	823.2	681.5
no quad	0.028	0.009	494.0	406.1	335.2	0.022	0.007	991.3	816.9	683.8
correct noise	0.031	0.010	493.9	402.7	334.8	0.022	0.007	991.4	816.4	677.3
correct prior	0.028	0.010	494.3	405.0	336.1	0.021	0.007	991.9	817.8	680.7
correct strong prior	0.027	0.010	494.2	404.7	337.2	0.021	0.008	992.6	821.0	677.3
unadjusted	0.026	0.016	494.6	426.0	367.0	0.024	0.010	990.2	859.6	734.2

adaptive simulations used in RCTs (United States Food and Drug Administration, 2019), our goal here is to compare model adjustment performance, not to obtain precise estimates of operating characteristics. The probability of stopping early is estimated as the proportion of times the trial stops before performing a final analysis. In the continuous and binary outcomes, this is the proportion of times the trial stops before enrolling the maximum number of participants. In the time-to-event outcome, this is the proportion of times the trial stops before 50 time units. In Bayesian adaptive designs which allow for early stopping, sample size is a random variable. Thus, expected sample size is estimated as the average of the 1,000 sample sizes at trial end. Posterior median bias is defined as the bias resulting from using the posterior median $\hat{\gamma}$ obtained from an adjustment model as an estimator for the value of γ used in the simulation, and is estimated through Monte Carlo using the 1,000 datasets. The Monte Carlo distribution of RMSE is displayed for each value of the marginal estimand. Here the entire posterior distribution from an adjustment model is used as the estimator for γ, so this is equivalent to the posterior expected squared error loss. For each of the 1,000

simulations, a single value of RMSE is obtained using the 3,000 posterior draws γ_s for the value of γ used in the simulation:

$$\text{RMSE} = \sqrt{\frac{1}{3000}\sum_{s=1}^{3000}(\gamma_s - \gamma)^2}.$$

Results for the continuous, binary, and time-to-event endpoints are are displayed in Figures 3.1-3.3 and Tables 3.2-3.4. For all endpoints, adjusting for variables known to be associated with the outcome increases the power of the trial and the probability of stopping the trial early as compared to the unadjusted analysis (Figures 3.1-3.3). Additionally, failing to correctly specify the functional form of a covariate (*no quad*) has only a minor impact on power and the probability of stopping early. Under all scenarios, incorporating stronger prior information appears to provide little to no benefit as compared to the weakly informative priors used in the *correct* models. This results from the priors being dominated by the data due to the high effective sample sizes. For the binary endpoint, this is induced by the control event risk of 0.3, which ensures that a moderately large number of events occurs throughout the trial. For the time-to-event endpoint, this is induced by the exponential baseline hazard rate of $\lambda = 0.02$, which ensures there are a large number of events within the maximum time limit of 50 time units.

Compared to other adjustment models, adjusting for noise (*correct noise*) has minimal impact under most scenarios. However, for the smallest maximum sample size under the binary endpoint (max ss = 100), adjustment for noise slightly increases power and the probability of stopping early as compared to the *correct* model. This may result from non-negligible correlation being induced between the outcome and noise variables under this setting, and so adjustment provides a further power benefit. However, this comes at the cost of a strong inflation in the Type 1 error rate as compared to the *correct* model (i.e., T1E = 0.079 versus T1E = 0.063 in Table 3.3). This underscores the importance of adjusting only for variables which are known to be associated with the outcome and in a pre-specified manner (Hauck

et al., 1998; Committee for Proprietary Medicinal Products, 2004; Senn, 2013; European Medicines Agency, 2015).

There is some suggestion that adjusted analyses tend to have slightly lower RMSE than unadjusted analyses under all scenarios. However, no clear pattern emerges for posterior median bias for the non-null treatment effects, where all adjustment models are comparable for practical purposes (Figures A.4.1-A.4.3 in Appendix A). For all scenarios except the smallest maximum sample size under the binary endpoint (max ss = 100), posterior median bias for the non-null treatment effects is negative. This results from the estimated treatment effect being further from the null than the true treatment effect (i.e., overestimation), which is expected in trials which allow for early stopping for treatment superiority (i.e., truncated trials) (Mueller et al., 2007; Walter et al., 2019; Robertson et al., 2023). For the null treatment effect, early stopping for superiority leads to non-zero but minimal bias under all endpoints and maximum sample sizes (Tables 3.2-3.4). For the binary endpoint, it is consistently larger in magnitude for the *unadjusted* model. We note that bias under the null is negative for the continuous endpoint but positive for the binary and time-to-event endpoints. This results from the marginal difference in means being unbounded below, whereas the marginal relative risk and marginal hazard ratio are bounded below by zero. When bias under the null is evaluated for these latter estimands on the log scale, most values become negative as in the continuous endpoint case. We elaborate further on this overestimation induced bias in Section A.6 of Appendix A.

Under all scenarios except the smallest maximum sample size (max ss = 100) for the binary endpoint, the Type 1 error rate is maintained below 0.05 for all adjustment models (Table 3.2-3.4). Under the smallest maximum sample size within the binary endpoint, however, all adjustment models lead to increased Type 1 error rate as compared to the *unadjusted* model. This results from using too many covariates in the adjustment model given the low effective sample size, a phenomenon known as over-stratification (Kahan et al., 2014). We observe

that as the maximum sample size, and thus effective sample size, increases, the inflation in the Type 1 error rate disappears. Under the time-to-event endpoint, there is minimal inflation in the Type 1 error rate for adjusted analyses as compared to the unadjusted analysis. This may also result from over-stratification as described above. Future work should determine the optimal number of covariates to include in an adjustment model under these scenarios to avoid over-stratification. Considering the substantial power gains achieved by adjusting under the time-to-event endpoint, and that the Type 1 error rate is maintained by selecting a conservative value of the probability of superiority threshold u, this slight increase in the Type 1 error rate is not likely to be problematic, however. Across all non-null treatment effect scenarios, the adjusted models have lower expected sample sizes than the unadjusted model. When combined with the probability of stopping early results, this implies that adjusted analyses are stopping more often and at earlier interim analyses for all endpoints. We note that the reduction in expected sample size is not as great for the time-to-event endpoint as compared to the other endpoints. This results from the maximum sample size being included for any interim analyses conducted past the halfway point of the trial, since all trial participants are enrolled by this point. A final simulation (included in Section A.5 of Appendix A) which incorporated varying degrees of prior information on the treatment effect was performed for the binary endpoint. This resulted in increased power and probability of stopping early for smaller maximum sample sizes, but at the cost of inflated type 1 error. A more complete investigation of including informative priors on treatment effects remains as future work.

3.5 Application: CCEDRRN-ADAPT

In this section we consider the design of a hypothetical platform trial which seeks to study the effectiveness of oral therapies against mild to moderate COVID-19 infection in individuals discharged from Canadian Emergency Departments. The trial design takes advantage of an already established network of physicians and researchers called the Canadian COVID-19

Emergency Department Rapid Response Network (CCEDRRN). We consider the first stage only, where a single oral therapy is compared to the standard of care. The binary outcome of interest is a composite endpoint of 28-day hospitalization or mortality. Realistic values used in the trial simulation performed below are taken from a COVID-19 Emergency Department risk model, developed by the CCEDRRN researchers (Hohl et al., 2022a). This risk model was developed using data from a high quality, population-based registry, which was also developed by the CCEDRRN researchers (Hohl et al., 2021,0). While the binary outcome for the risk model is all cause emergency department and in-hospital mortality, this is very likely to be highly correlated with the trial's composite outcome. Thus, the simulation's results are expected to generalize to the composite endpoint as well.

The data generating mechanism for the trial simulation is shown below, where a single maximum sample size of 3,000 is chosen due to the very low marginal control event risk ($p_{ctr} = 0.07$) and reflects the sample size used in CCEDRRN-ADAPT. Letting $Y = 1$ represent 28-day hospital admission or mortality, the marginal estimand of interest is the relative risk, $\gamma(\boldsymbol{\theta}) = \mu(\boldsymbol{\theta}; A = 1)/\mu(\boldsymbol{\theta}; A = 0)$ where $\mu(\boldsymbol{\theta}; A = a) = E[Y|A = a; \boldsymbol{\theta}]$. We assume $\gamma(\boldsymbol{\theta})$ is intention-to-treat, that is, we compare the outcome distributions of those who are assigned to treatment versus those who are assigned to control. As above, let η represent the linear predictor used in the data generating mechanisms. The set $\{\boldsymbol{\beta}, \phi\}$ represents the conditional covariate and treatment effects on the log-odds scale. As in the binary simulation case, the value of β_0 is optimized to ensure the generated datasets exhibit the correct marginal control event risk of $p_{ctr} = 0.07$. The values of ϕ are selected as those where the unadjusted model achieves approximately 50% and 80% power. The value of $\gamma = \gamma(\boldsymbol{\theta})$ corresponding to a specific value of ϕ is determined through simulation, described in Section A.2 of Appendix A. Thus, the values of ϕ and γ are reported together. Let F be a truncated normal distribution, parameterized by *post-truncation* values of its minimum $\wedge$, maximum $\vee$, first quartile $q1$, third quartile $q3$, mean μ, and standard deviation σ. An algorithm to simulate from this distribution is provided in Section A.3 of Appendix A.

The covariate distributions are simulated using values of summary statistics for the derivation cohort from a previously reported risk model for individuals presenting to Canadian Emergency Departments with COVID-19 symptoms (Hohl et al., 2022a). Covariates included in the risk model and whose summary statistics were available are included in the simulation. These include age in years (X_1), respiratory rate upon arrival to the Emergency Department and measured in breaths/minute (X_2), female sex (X_3), chest pain (X_4), and arrival from police or ambulance (X_5). As the risk model was developed before COVID-19 vaccines were widely available, vaccination status is not included as a potential covariate. The corresponding risk model coefficients are used as the conditional effects $\{\beta_1, ..., \beta_5\}$ in the simulation. Since the summary statistics did not include covariance information, joint independence between all variables is assumed. Due to study inclusion criteria, the range of the distribution of age is set to be $[18, 90]$. Due to biological contraints, the range of the distribution of respiratory rate is set to be $[12, 40]$. Values for age and respiratory rate are simulated from the truncated normal distribution F described above.

A total of 1,000 treatment-covariate datasets are generated, each containing 3,000 participants. Using these same 1,000 datasets, different outcome vectors are generated for each value of the marginal relative risk $\gamma \in \{1, 0.73, 0.63\}$, which corresponds to $\phi \in \{0, -0.42, -0.60\}$. The outcome vectors are generated as follows:

$$
\begin{aligned}
&Y \sim \text{Bernoulli}(\text{logit}^{-1}(\eta)) \\
&\eta = \beta_0 + \phi A + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_4 + \beta_5 X_5 && X_3 \sim \text{Bernoulli}(0.478) \\
&X_1 \sim F(\wedge = 18, \vee = 90, q1 = 39, q3 = 70, \mu = 54.7, \sigma = 19.8) && X_4 \sim \text{Bernoulli}(0.216) \\
&X_2 \sim F(\wedge = 12, \vee = 40, q1 = 18, q3 = 22, \mu = 21, \sigma = 6.2) && X_5 \sim \text{Bernoulli}(0.403) \\
&p_{ctr} = 0.07 && \Theta = \{\gamma \times \boldsymbol{\beta}\}. \\
&\boldsymbol{\beta} = (-10.76, 0.092, 0.097, -0.61, -0.80, 0.63)
\end{aligned}
$$

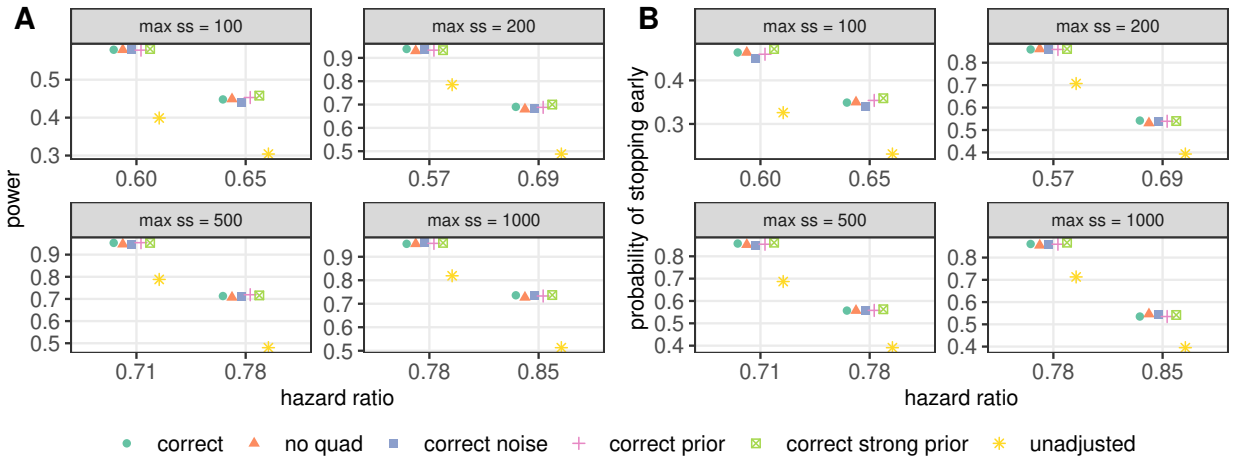

Figure 3.3: Time-to-event outcome. A) Power and B) probability of stopping early. Panels correspond to various maximum sample sizes (max ss). Points are jittered horizontally.

Table 3.4: Time-to-event outcome. Type 1 error rate (T1E), bias under the null (Bias*), and expected sample size at three different values of the marginal hazard ratio (γ).

	Maximum sample size = 100					Maximum sample size = 200				
			Expected sample size					Expected sample size		
Adjustment model	T1E	Bias*	$\gamma = 1$	$\gamma = 0.65$	$\gamma = 0.60$	T1E	Bias*	$\gamma = 1$	$\gamma = 0.69$	$\gamma = 0.57$
correct	0.027	0.014	99.7	98.3	98.0	0.033	0.000	199.4	193.1	189.4
no quad	0.027	0.017	99.7	98.3	97.9	0.029	0.002	199.4	193.5	189.7
correct noise	0.036	0.013	99.6	98.1	97.8	0.030	-0.001	199.4	193.1	189.2
correct prior	0.029	0.013	99.6	98.3	97.9	0.031	0.000	199.4	193.1	189.3
correct strong prior	0.026	0.015	99.7	98.3	97.8	0.033	0.000	199.3	193.3	189.2
unadjusted	0.029	0.035	99.7	98.5	98.2	0.028	0.005	199.4	194.4	191.0
	Maximum sample size = 500					Maximum sample size = 1000				
			Expected sample size					Expected sample size		
Adjustment model	T1E	Bias*	$\gamma = 1$	$\gamma = 0.78$	$\gamma = 0.71$	T1E	Bias*	$\gamma = 1$	$\gamma = 0.85$	$\gamma = 0.78$
correct	0.026	0.003	498.4	477.9	464.8	0.026	0.001	996.9	955.4	912.4
no quad	0.025	0.003	498.6	478.6	463.5	0.027	0.001	997.2	953.8	913.1
correct noise	0.027	0.003	498.6	477.6	464.2	0.022	0.001	997.4	953.3	910.7
correct prior	0.027	0.003	498.5	478.1	463.6	0.024	0.001	997.4	954.4	912.9
correct strong prior	0.026	0.003	498.4	477.8	463.2	0.026	0.001	997.4	955.4	911.9
unadjusted	0.022	0.011	498.6	483.4	470.6	0.019	0.003	997.9	962.8	929.6

The trial uses the same null and alternative hypotheses as those specified under the binary endpoint simulation above and includes a probability of superiority threshold of $u = 0.99$. Estimation for the marginal relative risk proceeds as previously described. Four adjustment models are considered and mirror those used in the binary simulations: *correct*, *correct prior*, *correct strong prior*, and *unadjusted*. All outcomes are assumed to be observed immediately upon participant enrollment. Interim analyses are event driven, and 75 new events are required to be observed before performing an interim analysis. This ensures a moderately high probability of stopping at an earlier interim analysis given treatment superiority. Interim analyses continue until the trial is stopped early for superiority or until 3,000 participants are enrolled, at which time the final analysis is performed.

Results for the CEDRRN-ADAPT design simulation study are summarized in Figure 3.4 and Table 3.5. Adjusting for variables known to be associated with the outcome increases the power of the trial and the probability of stopping the trial early as compared to the unadjusted analysis (Figure 3.4). As in the binary simulations above, including stronger priors on the covariate effects has minimal impact on power and the probability of stopping early as compared to the weakly informative priors used in the *correct* models. While there is some suggestion that adjusted analyses tend to have slightly lower RMSE than unadjusted analyses, all adjustment models have comparable posterior median bias for the non-null treatment effects (Figure A.4.4 in Appendix A). The Type 1 error rate is maintained below 0.05, and bias under the null is minimal, for all adjustment models (Table 3.5). As in the binary simulations above, there is slight inflation in bias under the null for the *unadjusted* model, however. The adjusted models have lower expected sample sizes than the unadjusted model across both non-null values of the relative risk. Again, we see the adjusted analyses are stopping more often and at earlier interim analyses as compared to the unadjusted analysis. For example, under a relative risk of 0.63, the *correct* model stops at the first interim analysis 50% of the time whereas the *unadjusted* model stops at the first interim analysis only 40% of the time.

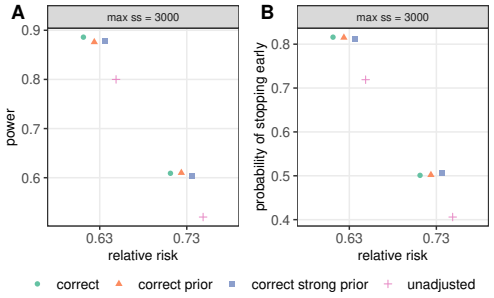

Figure 3.4: COVID-19 trial with binary outcome. A) Power and B) probability of stopping early. Points are jittered horizontally.

Table 3.5: COVID-19 trial with binary outcome. Type 1 error rate (T1E), bias under the null (Bias*), and expected sample size at three different values of the marginal relative risk (γ).

			Expected sample size		
Adjustment model	T1E	Bias*	$\gamma = 1$	$\gamma = 0.73$	$\gamma = 0.63$
correct	0.022	0.017	2969.7	2458.9	2047.3
correct prior	0.022	0.017	2968.9	2455.2	2051.8
correct strong prior	0.020	0.017	2971.5	2454.7	2056.0
unadjusted	0.024	0.022	2972.7	2561.7	2220.9

3.6 Discussion

The impact of covariate adjustment and incorporation of prior information on covariate effects has not been previously investigated in Bayesian adaptive trials with early stopping criteria. In this article, we assessed this impact using a variety of adjustment models and incorporated varying levels of prior information or types of model misspecification. It was shown that covariate adjustment increases power and the probability of stopping early, and decreases expected sample size over all scenarios. Furthermore, adjusting for covariates leads to trials which stop more often and at earlier interim analyses, and can decrease RMSE as compared to unadjusted analyses. These findings are fairly robust to adjustment for noise variables, but extra caution is needed for small sample size trials (max ss = 100) with binary endpoints where noise adjustment may lead to inflated Type 1 error. This reinforces existing best practice of only adjusting for covariates which have been pre-specified by subject matter experts (Hauck et al., 1998; Committee for Proprietary Medicinal Products, 2004; Senn, 2013; European Medicines Agency, 2015). For the scenarios explored here, which had moderate effective sample sizes, weakly informative priors on the covariate effects perform well, with stronger prior information providing little benefit. This included the CCEDRRN-ADAPT COVID-19 RCT design. Although the control event risk (0.07) of this trial is small, the maximum sample size was selected to be large enough to power the study for reasonable

values of relative risk, which ensured moderately high effective sample sizes. Including stronger prior information is expected to be helpful in trials which have small effective sample sizes (e.g., oncology trials), so future work may consider covariate adjustment within these contexts. Although we did not consider designs which also include futility stopping rules, we expect that the conclusions can be generalized.

In our simulation study, all covariates were assumed to be jointly independent. The assumption of independence may not hold in some cases and these should be investigated further. Since they carry similar information, adjusting for covariates which are moderately or highly correlated may yield smaller power increases than adjusting for approximately independent covariates. Future work might perform a simulation study to assess how different strengths of association between covariates impacts the results reported here. Another limitation of the current work is that adjustment was shown under only a single set of covariate effects within each endpoint. There is strong evidence in favor of covariate adjustment when the covariate effects are known to be strong. However, previous simulations (not shown) showed that adjustment under weaker covariate effects yielded only modest benefit. This finding was especially pronounced in the case of the non-linear endpoints. Future work might investigate a wider range of covariate effects to ascertain the magnitude that is required for covariate adjustment to be more than just modestly beneficial, though this may be context dependent.

It is important to note that the marginalization procedure used in this work is sensitive to the joint empirical distribution of the covariates with respect to which the conditional posterior samples are being marginalized. That is, the marginalization procedure is sensitive to the participants enrolled in the trial. While trialists work hard to obtain representative samples for use in RCTs, the participants volunteer and consent to be enrolled in the trials. Thus they are not a random sample from the population of interest (Lindsey and Lambert, 1998). If the sample is not as representative of the population of interest as desired, indirect standard-

ization could be performed. Here, conditional posterior samples would be marginalized with respect to a set of covariate patterns which more closely resembles the population of interest. Additionally, the participant covariate patterns could be augmented by *pseudo-participant* covariate patterns until the sample is more representative. These could be acquired from registry data or from participants in previous trials which contained a similar population and target of interest.

This work focused on Bayesian adaptive designs which employed simple randomization, where covariate adjustment takes place within an adaptive *decision* rule. Covariate adjustment may also occur within an adaptive *allocation* rule, such as in Covariate Adjusted Response Adaptive (CARA) designs (Rosenberger and Sverdlov, 2008; Villar and Rosenberger, 2018). These designs estimate treatment arm allocation probabilities from models which include covariates. While frequentist in nature, similar ideas can be applied within Bayesian response adaptive randomization designs. Since adaptive randomization leads to greater covariate imbalance across treatment arms, including covariate adjustment in both the adaptive decision and adaptive allocation rules may provide greater benefit than either do individually. This is an interesting direction for future work.

The simulation study above included variants of the *correct* adjustment model, which corresponded to the data generating mechanism used. In reality, this will be unknown. Variable selection and shrinkage methods might be employed to select covariates to be used in the adjustment model. Particularly, Bayesian model averaging can be used where the decision rules would use marginal posterior samples obtained from a consensus of plausible models and might be less sensitive to any single adjustment model misspecification. However, this approach is likely to be very computationally intense.

Another interesting direction of future research is exploration of Bayesian nonparametric models, such as Gaussian Processes, to consider more flexible functional forms that adjust for covariates. This approach might be especially advantageous in a setting where the underlying

association between the adjustment covariates and outcome is complex and hard to correctly specify, which might include a high degree of covariate non-linearity or interaction.

Data availability

Data sharing is not applicable to this article as no new data were created or analyzed in this study. The R scripts used for the simulations and graphics can be found on a public GitHub repository at https://github.com/jjwillard/cov_adj_bact.

Acknowledgements

The authors would like to thank Dr. Corinne Hohl for her friendly comments and suggestions regarding the CCEDRRN-ADAPT application component of this manuscript. JW acknowledges the support of a doctoral training scholarship from the Fonds de recherche du Québec - Nature et technologies (FRQNT). SG acknowledges the support by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada (NSERC) and a Chercheurs-boursiers (Junior 1) award from the Fonds de recherche du Québec - Santé (FRQS). EEMM acknowledges support from a Discovery Grant from NSERC. EEMM is a Canada Research Chair (Tier 1) in Statistical Methods for Precision Medicine and acknowledges the support of a chercheur de mérite career award from the FRQS. All simulations were performed on the Narval high performance computing cluster, located at the École de technologie supérieure in Montréal, Canada. JW thanks the research staff at Calcul Québec and the Digital Research Alliance of Canada (formerly Compute Canada) for their technical assistance.

Chapter 4

Bayesian optimization for personalized dose-finding trials with combination therapies

Preamble to Manuscript 2. I now shift focus to using covariate information in early phase clinical trials, by considering personalized dose-finding trials for combination therapies. Combination therapies combine two or more dosing agents. Personalized dose-finding seeks to identify optimal dose combinations based on individual patient characteristics. Many standard parametric dose-finding methods are difficult to extend to this setting since they potentially require many dose-covariate interaction terms, which may be poorly estimated given the small sample sizes in these trials. Thus, an important gap in the literature for personalized dose-finding trials for combination therapies remains. Mozgunov et al. (2022) considered dual-agent personalized dose-finding, but did not optimize the dose combination with respect to both dosing agents. Rather, the value of one of the agents for each patient was selected externally by clinicians based on patient-specific characteristics. The work in this chapter proposes using Bayesian optimization for multi-agent personalized dose-finding trials

where all dose dimensions are explored during the optimization and where early stopping is permitted if warranted by the data. The efficacy surface is modeled using a Gaussian process. After initial responses are collected, a sequential search policy is determined by optimizing an acquisition function. The performance of the procedure is assessed via a simulation study under multiple scenarios.

This work is motivated by an industry-sponsored problem, developed during an internship at a host organization, where interest is in the development of an intraocular implant combining two topical agents which have been in use for a long time and are well tolerated. Thus, no drug related AEs are expected and a minimal toxicity setting is assumed. Importantly, this means that toxicity information is ignored during the dose-finding procedure, which is viewed as an unconstrained optimization problem. Furthermore, response heterogeneity is expected to exist with respect to a key binary covariate, a characteristic of the eye, which motivates the personalized approach toward dose-finding. The methodological contributions of this manuscript are i) to propose a personalized dose-finding method for combination therapies which optimizes the dose combinations with respect to all dosing-agents, ii) to perform stratum-specific optimization by defining the acquisition function as a conditional expectation, and iii) to provide an early stopping rule which permits stratum-specific early stopping and reduces the expected sample size of the trial. The substantive contribution of this work is to demonstrate the feasibility of personalized dose-finding designs for combination therapies under a minimal toxicity setting despite the limited sample sizes used in early phase designs.

The corresponding manuscript has been submitted for review to a statistical journal.

Bayesian Optimization for Personalized Dose-Finding Trials with Combination Therapies

James Willard[1], Shirin Golchi[1], Erica E.M. Moodie[1], Bruno Boulanger[2], and Bradley P. Carlin[2]

[1]*Department of Epidemiology, Biostatistics, and Occupational Health, McGill University*

[2]*Cencora-PharmaLex*

Abstract

Identification of optimal dose combinations in early phase dose-finding trials is challenging, due to the trade-off between precisely estimating the many parameters required to flexibly model the possibly non-monotonic dose-response surface, and the small sample sizes in early phase trials. This difficulty is even more pertinent in the context of personalized dose-finding, where patient characteristics are used to identify tailored optimal dose combinations. To overcome these challenges, we propose the use of Bayesian optimization for finding optimal dose combinations in standard ("one size fits all") and personalized multi-agent dose-finding trials. Bayesian optimization is a method for estimating the global optima of expensive-to-evaluate objective functions. The objective function is approximated by a surrogate model, commonly a Gaussian process, paired with a sequential design strategy to select the next point via an acquisition function. This work is motivated by an industry-sponsored problem, where focus is on optimizing a dual-agent therapy in a setting featuring minimal toxicity. To compare the performance of the standard and personalized methods under this setting, simulation studies are performed for a variety of scenarios. Our study concludes that taking a personalized approach is highly beneficial in the presence of heterogeneity.

4.1 Introduction

Early phase clinical trials are designed to assess the safety and efficacy profiles of first in human doses of an experimental drug. Traditional adaptive dose-finding designs fall into three major families: algorithmic designs (e.g., Storer 1989), model-assisted designs (e.g., Yuan et al. 2016a), and model-based designs (e.g., O'Quigley et al. 1990). Here, our interest lies in model-based designs, which have been extended to the dual-agent dose combination setting (e.g., Wang and Ivanova 2005; Wages and Conaway 2014). Some of these designs (e.g., Houede et al. 2010; Wang et al. 2023) propose using flexible models to handle possible non-monotonicities in the dose-efficacy/dose-toxicity surfaces, recognizing that monotonicity depends on the type of drug being administered and may not hold in general (Li et al., 2017). Additionally, existing methods often restrict dose-finding to a small set of pre-selected dose combinations, which can fail to identify the optimal dose combination (Hirakawa et al., 2015). Adaptive dose insertion designs have been proposed to allow for the evaluation of additional dose combinations if warranted by the data (Cai et al., 2014; Lyu et al., 2019).

Most existing dose-finding designs, including those mentioned above, seek to find optimal dose combinations without consideration of additional covariate information (we refer to this as *standard dose-finding*). These designs ignore the possibility that patient responses may be variable across different subgroups within the population. With recent advances in molecular biology, specifically the identification of novel biomarkers, interest has grown in personalized (or precision) medicine. *Personalized dose-finding* aims to find optimal dose combinations based on individual patient characteristics. When utilizing parametric models, as in many of the standard dose-finding methods described above, extension to the personalized setting is challenging due to the limited sample sizes and potentially large number of dose-covariate interaction terms that are required to be estimated. Personalized dose-finding for monotherapies has been investigated using a variety of methods, including dimension-reduction techniques (e.g., Guo and Zang 2022), Bayesian hierarchical models (e.g., Morita

et al. 2017), and Bayesian model averaging (e.g. Psioda et al. 2021). The increased dimensionality in the combination therapy setting makes estimation even more challenging. Mozgunov et al. (2022) has considered personalized dose-finding for a dual-agent combination therapy where the patient-specific dose of one of the agents is selected externally by clinicians, but multi-agent personalized dose-finding designs where all dose dimensions are explored during optimization are still needed.

Our work is motivated by a problem in industry, where a sponsor is interested in a design for the development of an intraocular implant that combines two topical agents. Each agent has been in use and is well tolerated. No drug-related adverse events are expected, and we anticipate minimal toxicity, if any. Interest lies in obtaining the optimal dose combination of these two agents using a continuous efficacy measure, which need not be monotonic with the doses of each agent. Additionally, response heterogeneity is expected to exist with respect to a key binary covariate, a particular characteristic of the lens of the eye, and we seek a design which accommodates this feature. We allow for exploration of the dose combination region, which has been defined using data from previous observational and randomized studies, without consideration of formal dose escalation/de-escalation rules. As each new dose combination investigated carries engineering costs, early stopping is desirable if warranted by the data.

Considering this motivating example, we propose novel methodology for multi-agent dose-finding which meets several criteria. First, the model estimating the response surface is parsimonious, which is suitable for small sample sizes, yet remains flexible enough to model both monotonic and non-monotonic dose-response surfaces. Next, the design sequentially explores the entire dose combination space in a principled way, and allows for early stopping if warranted. Finally, our methodology facilitates extension to personalized dose-finding. Specifically, we propose modelling the response surface with a Gaussian process (GP) and using Bayesian optimization for efficient standard and personalized multi-agent dose-finding, as-

suming minimal toxicity. Bayesian optimization has been previously proposed for monotherapy in a standard dose-finding setting (Takahashi and Suzuki 2021a; Takahashi and Suzuki 2021b). Our contributions here differ from this previous work by utilizing a different sequential search policy, by extending the approach to combination therapies and personalized dose-finding setting, and by permitting early stopping.

The remainder of this manuscript is organized in the following manner. We first introduce Bayesian optimization, propose methods for standard and personalized multi-agent dose-finding, and describe an early stopping rule. A simulation study is then performed to compare the standard and personalized dose-finding methods for a dual-agent therapy under a variety of scenarios without early stopping. We then consider the development of the previously described intraocular implant and compare several dose-finding designs which include early stopping rules. We conclude with a discussion and possible directions for future work.

4.2 Bayesian optimization for dose-finding

Standard Dose-Finding. Let $f(\mathbf{d}) \in \mathbb{R}$ be a continuous dose-response surface of interest, defined as a function of dose combinations $\mathbf{d} = (d_1, ..., d_J)$ which are comprised of $J \in \mathbb{N}$ dosing agents and which lie in a continuous dose combination region $\mathbb{D} \subset \mathbb{R}^J$. In this setting, $f(\mathbf{d})$ may be the dose-efficacy surface or the dose-utility surface, where utility is defined using both efficacy and toxicity outcomes. For the remainder of the manuscript, we assume $f(\mathbf{d})$ is the dose-efficacy surface, which has been transformed such that *smaller* values are more desirable. We assume minimal toxicity over $\mathbb{D}$. Thus, dose escalation/de-escalation rules are not considered, and it is ethically permissible to select future dose combinations throughout the design space. Our interest is in obtaining the optimal dose combination $\mathbf{d}_{opt} \in \mathbb{D}$, defined as that which minimizes $f(\mathbf{d})$:

$$\mathbf{d}_{opt} = \underset{\mathbf{d} \in \mathbb{D}}{\operatorname{argmin}}\, f(\mathbf{d}).$$

To do so, we use Bayesian optimization, which is a derivative-free optimization method that estimates the global optima of expensive-to-evaluate objective functions (Garnett, 2023). It

is a sequential design strategy that approximates the objective function with a stochastic surrogate model (commonly a GP), and selects the next point by optimizing an *acquisition function*. Bayesian optimization is commonly used in engineering, machine learning, and computer experiments (Gramacy, 2020; Murphy, 2023), and has recently been employed to efficiently find optimal dynamic treatment regimes (Rodriguez Duque et al., 2022; Freeman et al., 2022). In this work, we model the efficacy function with a GP, and select the next dose combination as that which maximizes an acquisition function. This process repeats until optimality criteria are satisfied or sample size limits are reached.

To begin, a GP prior is placed on $f(\mathbf{d})$ (Williams and Rasmussen, 2006):

$$f(\mathbf{d}) \sim GP\left(m(\mathbf{d}), \nu\mathcal{K}(\mathbf{d}, \mathbf{d}')\right).$$

The GP is defined by its mean function $m(\mathbf{d})$ and correlation function (kernel) $\mathcal{K}(\mathbf{d}, \mathbf{d}')$, which is multiplied by scale parameter ν. The scale parameter determines the variability of the efficacy function throughout the dose combination space, and $\mathcal{K}(\mathbf{d}, \mathbf{d}')$ determines the correlation between responses for two dose combinations, $\mathbf{d}$ and $\mathbf{d}'$, where $\mathcal{K}(\mathbf{d}, \mathbf{d}') = 1$ when $\mathbf{d} = \mathbf{d}'$. Since the method's flexibility is largely determined by the choice of kernel function, constant mean GPs can be used, and in this work we set $m(\mathbf{d}) = \beta_0$. However, information about the response surface (e.g., pharmacokinetic/pharmacodynamic models) can be incorporated through the mean function if desired. Furthermore, we utilize a separable anisotropic squared exponential kernel,

$$\mathcal{K}(\mathbf{d}, \mathbf{d}') = \exp\left\{-\sum_{j=1}^{J} \frac{(d_j - d'_j)^2}{2l_{d_j}^2}\right\}. \tag{4.1}$$

This kernel is parameterized by characteristic length-scales l_{d_j}, which control the rate of decay in correlation between two dose combinations with respect to each dosing dimension and which are estimated using available data. Additionally, it is an example of a stationary kernel, which assumes that the degree of correlation between two dose combinations depends only on their distance and not on their locations in the dose combination space. We assume stationary

dose-response surfaces throughout this work, but comment on relaxing this assumption in the discussion.

To optimize $f(\mathbf{d})$, r patients are assigned to each of c initial dose combinations, yielding $n = r \times c$ patient responses which are treated as noisy observations $y_i = f(\mathbf{d}_i) + \epsilon_i$, where $\epsilon_i \sim N(0, \sigma_y^2)$ for $i = 1, ..., n$. This yields the observed data $\mathcal{D} = \{(\mathbf{d}_i, y_i)\}_{i=1}^n$, where we denote the vector of observations by $\mathbf{y} = (y_1, ..., y_n)^T$. Specifying the GP prior above induces a multivariate normal distribution on the observations, $\mathbf{y} \sim N(\beta_0 \mathbf{1}_n, \nu \mathbf{K})$, where $\mathbf{K}(i, j) = \mathcal{K}(\mathbf{d}_i, \mathbf{d}_j) + \tau^2 \mathbb{1}_{i=j}$ and τ^2 is a noise parameter. After observing $\mathcal{D}$, we adopt an empirical Bayes approach toward the kernel hyperparameters $\boldsymbol{\theta} = \{\nu, \tau^2, l_{d_1}, l_{d_2}\}$, replacing them with their maximum likelihood estimates (Gramacy, 2020). The posterior distribution of the efficacy function at a new dose combination, denoted by $\widetilde{\mathbf{d}}$, is then $p(f \mid \mathcal{D}, \widetilde{\mathbf{d}}) = N(\mu(\widetilde{\mathbf{d}}), \sigma^2(\widetilde{\mathbf{d}}))$ (Binois and Gramacy, 2021), such that

$$\begin{aligned} \mu(\widetilde{\mathbf{d}}) &= \hat{\beta}_0 + \mathbf{k}(\widetilde{\mathbf{d}})^T \mathbf{K}^{-1}(\mathbf{y} - \hat{\beta}_0 \mathbf{1}) \\ \sigma^2(\widetilde{\mathbf{d}}) &= \nu \mathcal{K}(\widetilde{\mathbf{d}}, \widetilde{\mathbf{d}}) - \nu \mathbf{k}(\widetilde{\mathbf{d}})^T \mathbf{K}^{-1} \mathbf{k}(\widetilde{\mathbf{d}}) + \frac{(1 - \mathbf{k}(\widetilde{\mathbf{d}})^T \mathbf{K}^{-1} \mathbf{1})^2}{\mathbf{1}^T \mathbf{K}^{-1} \mathbf{1}} \\ \hat{\beta}_0 &= \frac{\mathbf{1}^T \mathbf{K}^{-1} \mathbf{y}}{\mathbf{1}^T \mathbf{K}^{-1} \mathbf{1}} \end{aligned} \tag{4.2}$$

where $\mathbf{K}$ is $n \times n$, $\mathbf{k}(\widetilde{\mathbf{d}}) = (\mathcal{K}(\mathbf{d}_1, \widetilde{\mathbf{d}}), ..., \mathcal{K}(\mathbf{d}_n, \widetilde{\mathbf{d}}))^T$ is $n \times 1$, and $\hat{\beta}_0$ is a plug-in estimate for the mean.

The next dose combination for evaluation, denoted by $\mathbf{d}^{(c+1)}$, is then selected as that candidate dose combination $\widetilde{\mathbf{d}} \in \mathbb{D}$ which maximizes an acquisition function, denoted by $\alpha(\widetilde{\mathbf{d}} \mid \mathcal{D})$:

$$\mathbf{d}^{(c+1)} = \underset{\widetilde{\mathbf{d}} \in \mathbb{D}}{\operatorname{argmax}}\, \alpha(\widetilde{\mathbf{d}} \mid \mathcal{D}).$$

One acquisition function commonly paired with a GP is the Expected Improvement (EI), defined as

$$\alpha_{EI}(\widetilde{\mathbf{d}} \mid \mathcal{D}) = \mathbb{E}[\max(0, f^* - f(\widetilde{\mathbf{d}})) \mid \mathcal{D}, \widetilde{\mathbf{d}}]$$

and available in closed form (Jones et al., 1998),

$$\alpha_{EI}(\widetilde{\mathbf{d}} \mid \mathcal{D}) = (f^* - \mu(\widetilde{\mathbf{d}}))\Phi\left(\frac{f^* - \mu(\widetilde{\mathbf{d}})}{\sigma(\widetilde{\mathbf{d}})}\right) + \sigma(\widetilde{\mathbf{d}})\phi\left(\frac{f^* - \mu(\widetilde{\mathbf{d}})}{\sigma(\widetilde{\mathbf{d}})}\right),$$

where $f^* = \min_{\mathbf{d}_i} f(\mathbf{d}_i)$ denotes the value of the current observed optimum, $\Phi(\cdot)$ and $\phi(\cdot)$ denote the cdf and pdf of a standard normal random variable, respectively, and where $\mu(\widetilde{\mathbf{d}})$ and $\sigma(\widetilde{\mathbf{d}})$ are the posterior mean and standard deviation of the efficacy function evaluated at $\widetilde{\mathbf{d}}$, respectively. The EI balances between *exploiting* regions that have desirable values of $f(\mathbf{d})$ (first term), and *exploring* regions in the dose combination space that are imprecisely estimated (second term). Under a noisy setting, f^* is not observed and so plug-in estimates have been proposed, for example, by using the minimum of the posterior mean, $\widehat{f^*} = \min_{\widetilde{\mathbf{d}} \in \mathbb{D}} \mu(\widetilde{\mathbf{d}})$ (Gramacy and Lee, 2011; Picheny et al., 2013a). A different acquisition function, called the Augmented Expected Improvement (AEI; Huang et al. (2006)), has been shown to offer better performance than EI under higher noise settings (Picheny et al., 2013b). The AEI defines $\widehat{f^*} = \mu(\mathbf{d}^*)$, the posterior mean of f at the current "effective best solution", $\mathbf{d}^*$, which can be defined as the point which minimizes a posterior β-quantile. We follow the recommendation in Huang et al. (2006) and define $\mathbf{d}^*$ as that point which minimizes the $\beta = 0.84$ posterior quantile, which is equivalent to setting $\mathbf{d}^* = \text{argmin}_{\widetilde{\mathbf{d}}}\, \mu(\widetilde{\mathbf{d}}) + \gamma\sigma(\widetilde{\mathbf{d}})$ when $\gamma = 1$. The AEI is then defined as:

$$\alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D}) = \alpha_{EI}(\widetilde{\mathbf{d}} \mid \mathcal{D})\left(1 - \frac{\tau}{\sqrt{\sigma^2(\widetilde{\mathbf{d}}) + \tau^2}}\right). \tag{4.3}$$

The multiplicative term serves to promote exploration by penalizing dose combinations which have small posterior variance (Picheny et al., 2013b). In this work, we use AEI since we find it offers moderate improvement over EI.

After evaluation of $\mathbf{d}^{(c+1)}$ in r new patients, the data are updated, $\mathcal{D} = \mathcal{D} \cup \{(\mathbf{d}_i^{(c+1)}, y_i)\}_{i=1}^r$. The GP model is refit to obtain a new posterior distribution $p(\mathbf{f} \mid \mathcal{D}, \widetilde{\mathbf{d}})$. Then $s = 1, ..., S$ samples $f_s(\widetilde{\mathbf{d}})$ from this posterior are obtained to yield S samples from the posterior of the

optimal dose combination $p(\mathbf{d}_{opt} \mid \mathcal{D})$ as:

$$\mathbf{d}_{opt,s} = \underset{\widetilde{\mathbf{d}} \in \mathbb{D}}{\operatorname{argmin}}\, f_s(\widetilde{\mathbf{d}}).$$

This procedure continues until the sample size limit is reached or an early stopping rule, which we denote by $\mathbb{1}_{\text{STOP}}$, is satisfied.

One possible early stopping rule proposed in Huang et al. (2006) is to allow stopping only after $\max_{\widetilde{\mathbf{d}} \in \mathbb{D}} \alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D}) < \delta$. In this case, the algorithm terminates only if there is little improvement to be gained over f^* across the dose combination space. Under a noisy setting, the authors suggest this be satisfied for $(J+1)$ consecutive algorithm iterations before termination. This yields the following stopping rule:

$$\mathbb{1}_{\text{STOP}} = \begin{cases} \text{TRUE} & \max_{\widetilde{\mathbf{d}} \in \mathbb{D}} \alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D}) < \delta \text{ for } (J+1) \text{ iterations} \\ \text{FALSE} & \text{otherwise.} \end{cases} \tag{4.4}$$

We note that δ is a tuning parameter that controls performance of the algorithms. Its value can be determined through sensitivity analysis using several values obtained through Monte Carlo simulation. For example, the Monte Carlo distributions of α_{AEI} can be obtained at each iteration and different values of δ can be selected as summary statistics of these distributions (e.g., the median). The performance of these values can then be compared, with smaller values of δ implying more stringent stopping criteria, and larger values of δ permitting earlier stopping.

The sequential procedure and stopping rule described above suggest Algorithm 1, referred to as the *standard optimization algorithm*. In many applications of Bayesian optimization, it is standard to collect the initial data and then continue with $r = 1$ at each iteration of the algorithm. This permits maximal exploration of the domain, as a novel input point from the design space can be evaluated each time. In certain dose-finding applications, such as the motivating example of this work, engineering costs may prohibit the creation of a novel dose combination for each new patient, and some patients may necessarily be assigned to

Algorithm 1 Standard Optimization Algorithm

Input: r patient responses at each of c initial dose combinations, $\mathcal{D} = \{(\mathbf{d}_i, y_i)\}_{i=1}^{n_0}$

$n \leftarrow n_0 = r \times c$
$\mathbb{1}_{\text{STOP}} \leftarrow$ FALSE
Obtain $p(\mathbf{f} \mid \mathcal{D}, \widetilde{\mathbf{d}})$ and $p(\mathbf{d}_{opt} \mid \mathcal{D})$ using fitted GP ▷ Obtain posteriors
$\mathbf{d}^* \leftarrow \text{argmin}_{\widetilde{\mathbf{d}}}\, \mu(\widetilde{\mathbf{d}}) + \sigma(\widetilde{\mathbf{d}})$ ▷ Obtain effective best point
$f^* \leftarrow \mu(\mathbf{d}^*)$ ▷ Obtain best value of f
Calculate $\alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D})$ for $\widetilde{\mathbf{d}} \in \mathbb{D}$ ▷ Compute AEI
while $n < N$ and $\mathbb{1}_{\text{STOP}} =$ FALSE **do**
    $\mathbf{d}^{(c+1)} \leftarrow \text{argmax}_{\widetilde{\mathbf{d}} \in \mathbb{D}}\, \alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D})$ ▷ Obtain next dose
    for $i = 1, ..., r$ **do**
        Evaluate y_i at $\mathbf{d}^{(c+1)}$ ▷ Observe outcomes
    end for
    $n \leftarrow n + r$; $c \leftarrow c + 1$ ▷ Update n, c
    $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\mathbf{d}_i^{(c+1)}, y_i)\}_{i=1}^{r}$ ▷ Update data
    Obtain $p(\mathbf{f} \mid \mathcal{D}, \widetilde{\mathbf{d}})$ and $p(\mathbf{d}_{opt} \mid \mathcal{D})$ using fitted GP ▷ Obtain posteriors
    $\mathbf{d}^* \leftarrow \text{argmin}_{\widetilde{\mathbf{d}}}\, \mu(\widetilde{\mathbf{d}}) + \sigma(\widetilde{\mathbf{d}})$ ▷ Obtain effective best point
    $f^* \leftarrow \mu(\mathbf{d}^*)$ ▷ Obtain best value of f
    Calculate $\alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D})$ for $\widetilde{\mathbf{d}} \in \mathbb{D}$ ▷ Compute AEI
    Update $\mathbb{1}_{\text{STOP}}$ using (4.4) ▷ Update stopping rule
end while

the same dose.

Personalized Dose-Finding. Personalized medicine recognizes that response heterogeneity may exist within the population of interest; personalized dose-finding incorporates covariate information to account for this. Consider a set of P discrete covariates $\mathbf{Z} = \{Z_p\}_{p=1}^{P}$. The Cartesian product of the levels of these P covariates define K strata. Personalized dose finding seeks to find the optimal dose combinations, denoted by $\mathbf{d}_{opt,k}$, across the continuous dose combination space for each of the K strata:

$$\mathbf{d}_{opt,k} = \underset{\mathbf{d} \in \mathbb{D}}{\text{argmin}}\, f(\mathbf{d}, \mathbf{z}_k) \quad \text{for} \quad k = 1, ..., K.$$

The efficacy function $f(\mathbf{d}, \mathbf{z})$ is modeled using a single GP fit to the data $\mathcal{D} = \{(\mathbf{d}_i, \mathbf{z}_i, y_i)\}_{i=1}^{n}$, where $y_i = f(\mathbf{d}_i, \mathbf{z}_i) + \epsilon_i$ with $\epsilon_i \sim N(0, \sigma_y^2)$, which allows information to be borrowed across strata. One possible method of incorporating the additional covariates into the GP model is through the kernel function. We use a (stationary) separable anisotropic squared exponential

kernel function that includes the additional covariates of interest:

$$\mathcal{K}((\mathbf{d},\mathbf{z}),(\mathbf{d},\mathbf{z})') = \exp\left\{-\left(\sum_{j=1}^{J}\frac{(d_j-d_j')^2}{2l_{d_j}^2}+\sum_{p=1}^{P}\frac{(z_p-z_p')^2}{2l_p^2}\right)\right\}. \tag{4.5}$$

As before, we use an empirical Bayes approach for estimating the hyperparameters. In the case where Z_p is a binary variable representing the levels of covariate p, the correlation between two patient responses is reduced by a factor of $\exp(-1/(2l_p^2))$ if they belong to different strata with respect to covariate p.

Since the efficacy function may exhibit different behavior within each stratum, each stratum may have a unique best efficacy function value f_k^*. Similar to the standard case, f_k^* is estimated as the posterior mean at the effective best point within stratum k, $f_k^* = \mu(\mathbf{d}^*, \mathbf{z}_k)$. To account for this heterogeneity, the sequential selection is performed within each stratum but with the GP fit utilizing data from all strata. This proceeds by modifying the AEI acquisition function to use f_k^* when conditioned on being in stratum k, denoted by $\alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D}, \widetilde{\mathbf{z}}_k)$. The sequential procedure continues until sample size limits are reached or until an early stopping rule is satisfied for each stratum. Since optimal doses in some strata may be easier to identify than in others, stratum-specific early stopping should be employed. One possible stratum specific stopping rule, which we denote by $\mathbb{1}_{\text{STOP},k}$, replaces $\alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D})$ in (4.4) with $\alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D}, \widetilde{\mathbf{z}}_k)$. Throughout this work, we assume that once dose-finding is stopped in a stratum, it is not re-started in that stratum. This, of course, can be changed in practice if desired. Upon termination of the algorithm, the posterior distribution of the optimal dose combination for each stratum is returned, denoted by $p(\mathbf{d}_{opt} \mid \mathcal{D}, \mathbf{z}_k)$, and uses the final GP model fitted on all of the data. These modifications suggest Algorithm 2, the *personalized optimization algorithm*.

4.3 Simulation study

Below we perform a simulation study to compare the performance of the standard and personalized algorithms (Algorithms 1 and 2, respectively) under three scenarios with no

Algorithm 2 Personalized Optimization Algorithm

Input: r patient responses at each of c_k initial dose combinations per stratum, $\mathcal{D} = \{(\mathbf{d}_i, \mathbf{z}_i, y_i)\}_{i=1}^{n_0}$

$n \leftarrow n_0 = \sum_{k=1}^{K} n_{0,k}$
Obtain $p(\mathbf{f} \mid \mathcal{D}, \widetilde{\mathbf{d}}, \widetilde{\mathbf{z}})$ using fitted GP ▷ Obtain posterior of f
for $k = 1, ..., K$ **do**
    Obtain $p(\mathbf{d}_{opt} \mid \mathcal{D}, \mathbf{z}_k)$ ▷ Obtain posterior of $\mathbf{d}_{opt,k}$
    $\mathbf{d}^* \leftarrow \text{argmin}_{\widetilde{\mathbf{d}}}\, \mu(\widetilde{\mathbf{d}}, \mathbf{z}_k) + \sigma(\widetilde{\mathbf{d}}, \mathbf{z}_k)$ ▷ Obtain effective best point
    $f_k^* \leftarrow \mu(\mathbf{d}^*, \mathbf{z}_k)$ ▷ Obtain best value f_k^*
    Calculate $\alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D}, \widetilde{\mathbf{z}}_k)$ for $\widetilde{\mathbf{d}} \in \mathbb{D}$ ▷ Compute AEI
    $\mathbb{1}_{\text{STOP},k} \leftarrow$ FALSE
end for
while $n < N$ and $\mathbb{1}_{\text{STOP},k} =$ FALSE for at least one stratum k **do**
    for $k = 1, ..., K$ **do**
        $n_k \leftarrow 0$
        $\mathcal{D}_k = \emptyset$
        if $\mathbb{1}_{\text{STOP},k} =$ FALSE **then**
            $\mathbf{d}^{(c_k+1)} \leftarrow \text{argmax}_{\tilde{\mathbf{d}} \in \mathbb{D}}\, \alpha_{AEI}(\tilde{\mathbf{d}} \mid \mathcal{D}, \widetilde{\mathbf{z}}_k)$ ▷ Obtain next dose
            for $i = 1, ..., r$ **do**
                Evaluate y_i at $(\mathbf{d}^{(c_k+1)}, \mathbf{z}_k)$ ▷ Observe outcomes
            end for
            $n_k \leftarrow r$; $c_k \leftarrow c_k + 1$ ▷ Update n_k, c_k
            $\mathcal{D}_k = \{(\mathbf{d}_i^{(c_k+1)}, \mathbf{z}_k, y_i)\}_{i=1}^{r}$
        end if
    end for
    $n \leftarrow n + \sum_{k=1}^{K} n_k$ ▷ Update n
    $\mathcal{D} = \mathcal{D} \bigcup_{k=1}^{K} \mathcal{D}_k$ ▷ Update data
    Obtain $p(\mathbf{f} \mid \mathcal{D}, \widetilde{\mathbf{d}}, \widetilde{\mathbf{z}})$ using fitted GP ▷ Obtain posterior of f
    for $k = 1, ..., K$ **do**
        Obtain $p(\mathbf{d}_{opt} \mid \mathcal{D}, \mathbf{z}_k)$ ▷ Obtain posterior of $\mathbf{d}_{opt,k}$
        if $\mathbb{1}_{\text{STOP},k} =$ FALSE **then**
            $\mathbf{d}^* \leftarrow \text{argmin}_{\widetilde{\mathbf{d}}}\, \mu(\widetilde{\mathbf{d}}, \mathbf{z}_k) + \sigma(\widetilde{\mathbf{d}}, \mathbf{z}_k)$ ▷ Obtain effective best point
            $f_k^* \leftarrow \mu(\mathbf{d}^*, \mathbf{z}_k)$ ▷ Obtain best value f_k^*
            Calculate $\alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D}, \widetilde{\mathbf{z}}_k)$ for $\widetilde{\mathbf{d}} \in \mathbb{D}$ ▷ Compute AEI
            Update $\mathbb{1}_{\text{STOP},k}$ using (4.4) ▷ Update stopping rule
        end if
    end for
end while

early stopping (i.e., $\delta = 0$ in (4.4)). Early stopping will be investigated in the next section. Scenarios 1 and 2 consider a single binary covariate Z_1. We index the true optimal dose combinations, $\mathbf{d}_{opt,k}$, and the true optimal values of the efficacy function, $f_{opt,k}$, using the

values of Z_1. That is, when $Z_1 = 0$ we use $\mathbf{d}_{opt,0}$ and $f_{opt,0}$. Scenario 3 considers two binary covariates, Z_1 and Z_2, and the $\mathbf{d}_{opt,k}$ and $f_{opt,k}$ are indexed similarly. For example, when $Z_1 = 0$ and $Z_2 = 1$, we use $\mathbf{d}_{opt,01}$ and $f_{opt,01}$. To make the simulations in this manuscript more computationally feasible, we modify Algorithms 1 and 2 to return a point estimate of $\mathbf{d}_{opt,k}$ rather than the entire posterior distribution. The point estimate is defined as the minimizer of the posterior mean surface, $\widehat{\mathbf{d}}_{opt,k} = \text{argmin}_{\widetilde{\mathbf{d}}}\, \mu(\widetilde{\mathbf{d}}, \widetilde{\mathbf{z}}_k)$.

We utilize dose combinations $\mathbf{d} = (d_1, d_2) \in [0,1]^2$, assumed to be standardized, where $\mathbf{d} = (0,0)$ corresponds to the combination using the lowest doses of interest for each agent and where $\mathbf{d} = (1,1)$ corresponds to the combination using the highest doses of interest for each agent. The point estimates, $\widehat{\mathbf{d}}_{opt,k}$, and the next dose combinations for evaluation, $\mathbf{d}^{(c_k+1)}$, are set, respectively, as the minimizers of $\mu(\widetilde{\mathbf{d}}, \widetilde{\mathbf{z}}_k)$ and maximizers of $\alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D}, \widetilde{\mathbf{z}}_k)$. The $\alpha_{AEI}(\widetilde{\mathbf{d}} \mid \mathcal{D}, \widetilde{\mathbf{z}}_k)$ is evaluated across an evenly spaced grid on $[0,1]^2$. The grid is incremented by 0.25 in each dimension, reflecting the degree of precision to which the drug maker can manufacture a particular dose combination. As a result, some dose combinations may be suggested more than once in the algorithm. While the proposed method is capable of optimizing over the continuous dose combination space, it is important to incorporate any manufacturing constraints into the optimization procedures to avoid suggesting doses which are not feasible to engineer. Additionally, the proposed method can be used for unevenly spaced grids on the dose combination space, though we do not consider that scenario here.

As before, we define $f(\mathbf{d}, \mathbf{z})$ to be a continuous efficacy surface and assume we are in a minimal toxicity setting where it is ethically permissible to select future doses anywhere in the dose combination space. Let g_i be a bivariate normal density function, $N(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$, which is parameterized by mean vector $\boldsymbol{\mu}_i$ and covariance matrix $\boldsymbol{\Sigma}_i$, for $i = 1, 2, 3$. The efficacy functions under the considered scenarios use the following densities evaluated at $\mathbf{d}$, denoted

Table 4.1: Simulation scenarios considered. The data generating mechanism for each scenario is $y = f(\mathbf{d}, \mathbf{z}) + \epsilon$ where $\epsilon \sim N(0, \sigma_y^2)$. The table columns contain the location of the optimal dose combination ($\mathbf{d}_{opt}$), the optimal value of the efficacy function (f_{opt}), the standardized effect size (ses), and whether or not the dose-efficacy surface is monotonically increasing with respect to each dosing dimension (monotone).

	$f(\mathbf{d}, \mathbf{z})$	σ_y	z_1	z_2	$\mathbf{d}_{opt}$	f_{opt}	ses	monotone
Simulation Study	1) $-\mathbb{1}\{z_1 = 0\} \times g_1(\mathbf{d}) -$	2.015	0		$(1, 1)$	-1.592	0.79	Yes
	$\mathbb{1}\{z_1 = 1\} \times g_1(\mathbf{d})$		1		$(1, 1)$	-1.592	0.79	Yes
	2) $-\mathbb{1}\{z_1 = 0\} \times g_2(\mathbf{d}) -$	0.319	0		$(0.25, 0.75)$	-1.203	3.77	No
	$\mathbb{1}\{z_1 = 1\} \times g_3(\mathbf{d})$		1		$(0.75, 0.25)$	-1.203	3.77	No
	3) $-\mathbb{1}\{(z_1 = 0, z_2 = 0)\} \times 0 -$	1	0	0	None	None	0	No
	$\mathbb{1}\{(z_1 = 0, z_2 = 1)\} \times 0.831 g_2(\mathbf{d}) -$		0	1	$(0.25, 0.75)$	-1	1	No
	$\mathbb{1}\{(z_1 = 1, z_2 = 0)\} \times 3.134 g_3(\mathbf{d}) -$		1	0	$(0.75, 0.25)$	-3.77	3.77	No
	$\mathbb{1}\{(z_1 = 1, z_2 = 1)\} \times 0.496 g_1(\mathbf{d})$		1	1	$(1, 1)$	-0.79	0.79	Yes
Implant	1) $-\mathbb{1}(z_1 = 0) \times 2.49 g_2(\mathbf{d}) -$	5	0		$(0.25, 0.75)$	-5	1	No
	$\mathbb{1}(z_1 = 1) \times 6.65 g_3(\mathbf{d}) - 2$		1		$(0.75, 0.25)$	-10	2	No

Note: The subtraction of 2 in $f(\mathbf{d}, \mathbf{z})$ under the Implant scenario corresponds to a base level of drug response outside the regions of optimality.

by $g_i(\mathbf{d})$ in Table 4.1:

$$g_1 = N\left[\begin{pmatrix}1\\1\end{pmatrix}, \begin{pmatrix}0.1 & 0\\0 & 0.1\end{pmatrix}\right] \quad g_2 = N\left[\begin{pmatrix}0.25\\0.75\end{pmatrix}, \begin{pmatrix}0.2 & 0.05\\0.05 & 0.1\end{pmatrix}\right] \quad g_3 = N\left[\begin{pmatrix}0.75\\0.25\end{pmatrix}, \begin{pmatrix}0.2 & 0.05\\0.05 & 0.1\end{pmatrix}\right].$$

The data generating mechanism for each scenario is $y = f(\mathbf{d}, \mathbf{z}) + \epsilon$ where $\epsilon \sim N(0, \sigma_y^2)$, with the specification of $f(\mathbf{d}, \mathbf{z})$ included in the first panel of Table 4.1 (rows labeled "Simulation Study") and plotted in panel A of Figures 4.1-4.3. The values of σ_y are chosen to ensure specific standardized effect sizes, defined as $ses = |f_{opt}|/\sigma_y$. We consider several standardized effect sizes drawn from a meta-analysis of dose-responses for a large drug development portfolio at a pharmaceutical company (Thomas et al., 2014). We focus on standardized effect sizes for drugs that had laboratory confirmed endpoints, which is the type of endpoint used in our motivating example. The $25^{th}/50^{th}/75^{th}$ percentiles of these standardized effect sizes are 0.79/1/3.77, which we refer to as small/medium/large effect sizes.

Scenario 1 considers the case of no response heterogeneity across a binary covariate Z_1 and

includes a small standardized effect size with a dose-efficacy surface which is monotonically increasing with respect to each dosing dimension. Both the locations of the optimal dose combinations, and the optimal values of the efficacy function, are the same across the strata. That is, $\mathbf{d}_{opt,0} = \mathbf{d}_{opt,1}$ and $f_{opt,0} = f_{opt,1}$. Scenario 2 considers response heterogeneity across Z_1, where the locations of the optimal dose combinations differ. Under this scenario, $\mathbf{d}_{opt,0} \neq \mathbf{d}_{opt,1}$ but $f_{opt,0} = f_{opt,1}$. This scenario considers large standardized effect sizes with efficacy surfaces which are non-monotone with respect to each dosing dimension. Scenario 3 considers heterogeneity across two binary covariates, Z_1 and Z_2, where both the locations of the optimal dose combinations and the optimal values of the efficacy function differ across the strata. Thus, $\mathbf{d}_{opt,ij} \neq \mathbf{d}_{opt,lm}$ for $ij \neq lm$ and $f_{opt,ij} \neq f_{opt,lm}$ for $ij \neq lm$. This scenario includes a zero stratum, $(z_1 = 0, z_2 = 0)$, which represents the covariate pattern of those who do not respond to the drug. We note that in this stratum, $\mathbf{d}_{opt,00}$ and $f_{opt,00}$ do not exist, but we consider the standardized effect size to be 0. This scenario includes small, medium, and large standardized effect sizes as well as both monotone and non-monotone dose-efficacy surfaces.

For each scenario, the standard and personalized algorithms are run using a maximum sample size of 80 participants. While this number is larger than many early phase trials might be in practice, our goal is to investigate the algorithms' performance characteristics as the sample size increases. We defer the use of early stopping rules to the following section, but note that these will permit a reduction in the expected sample size. We assign participants to dose combinations such that the total sample size of each algorithm is equal at each iteration, which allows a comparison of their performance. For the standard algorithm under scenarios 1 and 2, $r = 4$ participants are assigned to each dose combination on an initial dose matrix comprised of $c = 5$ dose combinations, which are selected via a two-dimensional quasi-random Sobol sequence (Morgan-Wall, 2022). This sequence serves as a space filling design and seeks to spread out the initial dose combinations in a more uniform manner than is typically accomplished via random sampling. More than one patient is assigned per dose combination

to control the cost associated with producing novel dose-combinations, a financial constraint from our motivating problem. This yields $n_0 = 20$. At each iteration of the algorithm, $r = 4$ additional participants are assigned to each proposed dose combination. This yields a total sample size of 80 after 15 iterations. We note that r is a tuning parameter, and that reducing it will lead to more unique doses being explored for the same fixed sample size. We explore this in the next section.

For the personalized algorithm, $c = 5$ initial dose combinations are selected in the same manner as above. To achieve the same total sample size by iteration as the standard algorithm, only $r = 2$ participants are assigned to each dose combination within each stratum. This yields $n_{0,0} = n_{0,1} = 10$, and so $n_0 = 20$, as in the standard algorithm. At each iteration within each stratum, $r = 2$ additional participants are assigned to each proposed dose combination. This yields a total sample size of 80 after 15 iterations. We note that this equal allocation of patients across strata represents an idealized trial which assumes that both the prevalence and enrollment of the specific subgroups is the same throughout the course of the trial. We consider departures from these assumptions in the discussion section. The same setup is used for scenario 3. However, since the number of strata is doubled, the number of participants evaluated at each dose combination is halved for the personalized algorithm. Thus, the standard algorithm still evaluates $r = 4$ participants per dose combination, but the personalized algorithm evaluates $r = 1$ participant per dose combination, yielding a total sample size of 80 after 15 iterations. All computing is performed in the statistical programming language R (R Core Team, 2022). The efficacy functions are modeled using constant mean GP models which utilize the anisotropic separable squared exponential kernels previously described, with the hyperparameters being jointly optimized by maximizing the marginal log-likelihood of the data (i.e., empirical Bayes GP; Binois and Gramacy 2021).

Algorithm performance is compared using several criteria which are estimated via $m =$

$1, ..., 1000$ Monte Carlo simulations. The expected number of dosing units from the optimal dose combination is used to assess how close the recommended dose combination $\widehat{\mathbf{d}}_k$ is to $\mathbf{d}_{opt,k}$ at each iteration. This measure is defined as the expected value of the Euclidean distance between $\widehat{\mathbf{d}}_k$ and $\mathbf{d}_{opt,k}$ divided by the precision to which the sponsor can manufacture doses, which is 0.25 in our simulations:

$$E_{y|z_k}[\text{dose units}] \approx \frac{1}{1000}\sum_{m=1}^{1000}\frac{\sqrt{(\hat{d}_1^{(m)} - d_{1,opt})_k^2 + (\hat{d}_2^{(m)} - d_{2,opt})_k^2}}{0.25}. \tag{4.6}$$

We utilize the average root posterior squared error loss (RPSEL) to assess how well the true efficacy function value $f(\widehat{\mathbf{d}}_k, z_k)$ at the recommended dose is estimated by the pointwise posterior distribution of the efficacy function at the recommended dose, $p(f_k \mid \mathcal{D}, \widehat{\mathbf{d}}_k, z_k)$, where $f^{(s)}(\widehat{\mathbf{d}}_k, z_k)$ denotes a single posterior sample out of $s = 1, ..., 10000$ posterior samples:

$$\text{Average RPSEL} \approx \frac{1}{1000}\sum_{m=1}^{1000}\left[\frac{1}{10000}\sum_{s=1}^{10000}(f^{(s)}(\widehat{\mathbf{d}}_k^{(m)}, z_k) - f(\widehat{\mathbf{d}}_k, z_k))^2\right]^{\frac{1}{2}}. \tag{4.7}$$

We employ $f(\widehat{\mathbf{d}}_k, z_k)$ here rather than $f_{opt,k}$ to understand how well the algorithms capture f at the recommended dose combination even if the recommended dose combination is not optimal. This is important for later phase studies which may utilize estimates of f obtained at the recommended dose combination for sample size and power calculations. In a similar manner, we present the average absolute deviation of the posterior mean point estimates $E[f_k \mid \mathcal{D}, \widehat{\mathbf{d}}_k, z_k]$ from the true value of $f(\widehat{\mathbf{d}}_k, z_k)$ at each iteration. Note that the standard algorithm ignores the strata and so yields only a single recommended dose $\widehat{\mathbf{d}}_k = \widehat{\mathbf{d}}$, posterior distribution $p(f_k \mid \mathcal{D}, \widehat{\mathbf{d}}_k, z_k) = p(f \mid \mathcal{D}, \widehat{\mathbf{d}})$, and posterior mean $E[f_k \mid \mathcal{D}, \widehat{\mathbf{d}}_k, z_k] = E[f \mid \mathcal{D}, \widehat{\mathbf{d}}]$ per iteration for $k = 1, ..., K$. See panels B-D of Figures 4.1-4.3 for these criteria by iteration.

Scenario 1 considers the case of no response heterogeneity across a single binary covariate Z_1 for small standardized effect sizes with monotonically increasing dose-efficacy surfaces. Under this scenario, both algorithms converge to the locations of $\mathbf{d}_{opt,k}$, coming within one

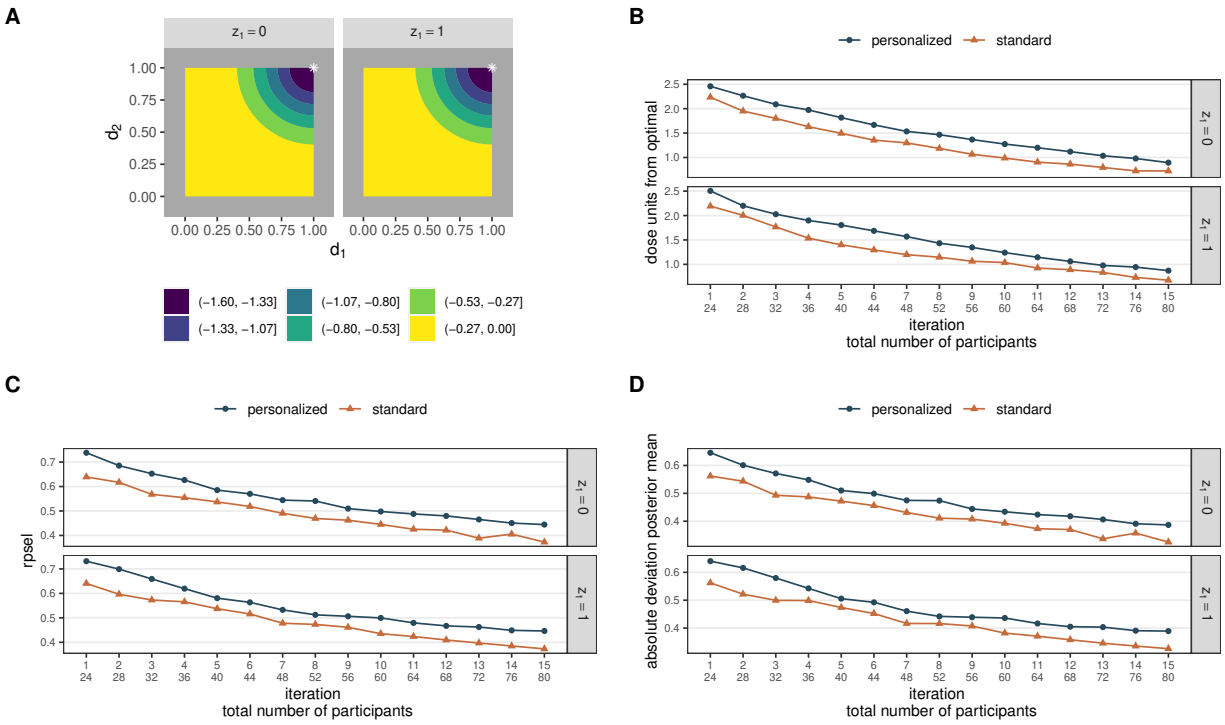

Figure 4.1: Scenario 1. A) Efficacy function, white stars denote $\mathbf{d}_{opt,k}$, B) expected dose units from the optimal dose combination as defined in (4.6) by iteration, C) average RPSEL as defined in (4.7) by iteration, and D) average absolute deviation of the posterior mean estimate by iteration.

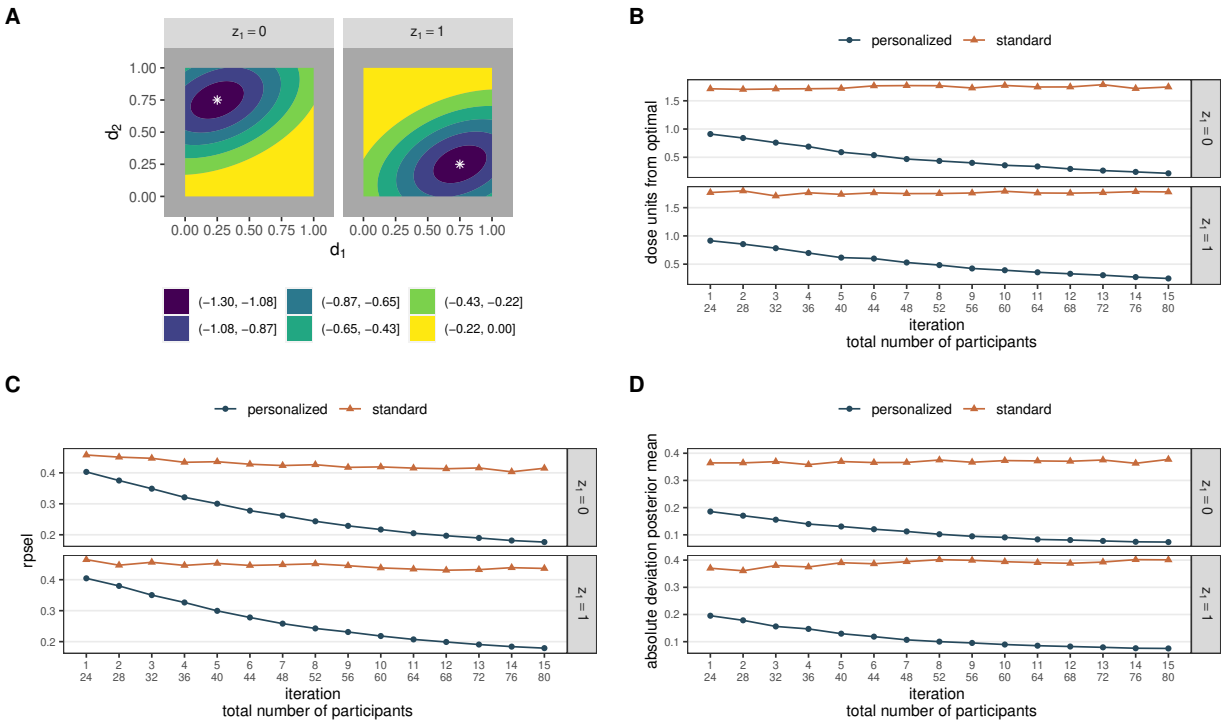

Figure 4.2: Scenario 2. A) Efficacy function, white stars denote $\mathbf{d}_{opt,k}$, B) expected dose units from the optimal dose combination as defined in (4.6) by iteration, C) average RPSEL as defined in (4.7) by iteration, and D) average absolute deviation of the posterior mean estimate by iteration.

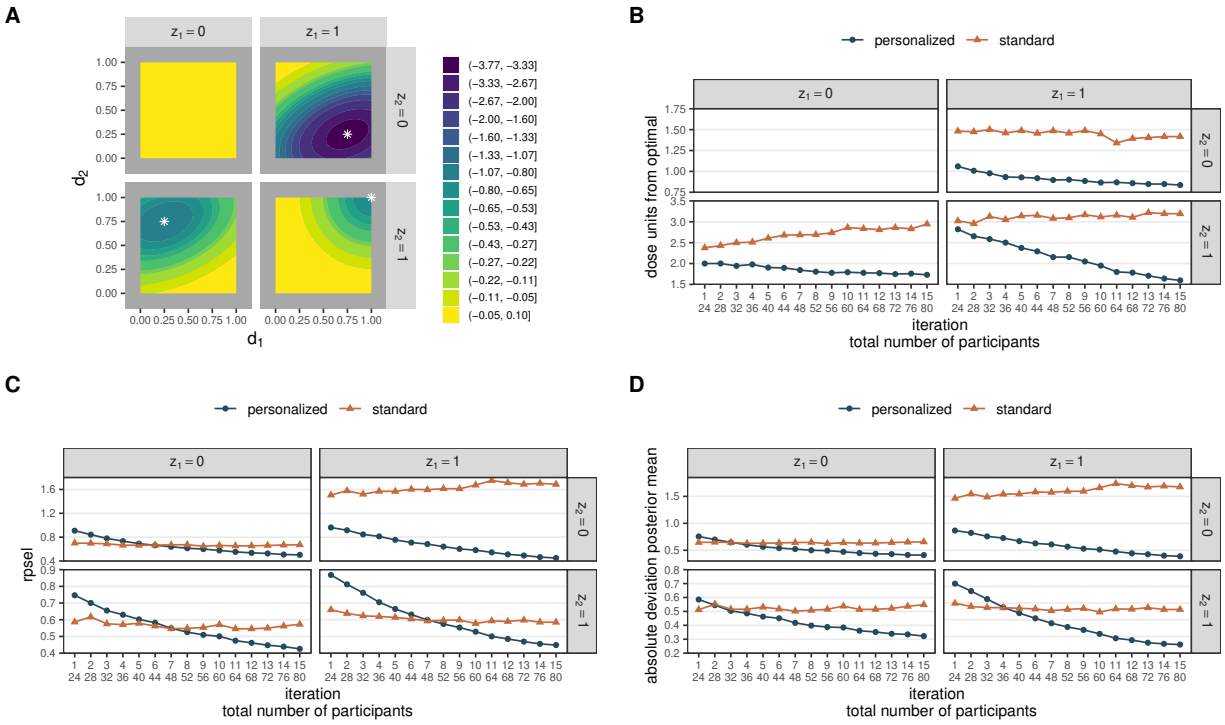

Figure 4.3: Scenario 3. A) Efficacy function, white stars denote $\mathbf{d}_{opt,k}$, B) expected dose units from the optimal dose combination as defined in (4.6) by iteration, C) average RPSEL as defined in (4.7) by iteration, and D) average absolute deviation of the posterior mean estimate by iteration. In (B), the plot for $(z_1 = 0, z_2 = 0)$ is empty as no optimal dose combination exists.

dosing unit of the optimal (Panel B of Figure 4.1). Estimation of f is challenging for the small standardized effect size (i.e., higher level of noise), but by termination, the algorithms come within 0.4 units of the true f at the recommended dose combination (panels C-D of Figure 4.1). The personalized algorithm is slightly less efficient than the standard algorithm, however, and takes longer to converge. This is likely the result of the GP model used in the personalized algorithm needing to estimate the additional length-scale parameter for Z_1, l_1.

Scenario 2 considers response heterogeneity across Z_1 for large standardized treatment effect sizes with non-monotonic dose-efficacy surfaces. Under this scenario, the personalized algorithm converges to the locations of $\mathbf{d}_{opt,k}$ and estimates the true value of f at the recommended dose combination well, whereas the standard algorithm does not (panels B-D of Figure 4.2). The standard algorithm typically explores the area near only a single $\mathbf{d}_{opt,k}$ or attempts to explore both optima (not shown). This results from the marginal efficacy function surface being bimodal, since it is a mixture distribution comprised of the equally weighted strata of Z_1, which are displayed in panel A of Figure 4.2. Note that even if the bimodality of the marginal surface is properly identified and explored, patients cannot be optimally treated without consideration of Z_1. That is, without this additional covariate information, the standard algorithm cannot determine which mode should be used to treat patients with $Z_1 = 0$ versus $Z_1 = 1$. This is possible using the personalized approach, however.

Scenario 3 considers heterogeneity across two binary covariates, Z_1 and Z_2, and includes zero, small, medium, and large standardized effect sizes with both monotone and non-monotone dose-efficacy surfaces. Recall that the stratum $(z_1 = 0, z_2 = 0)$ corresponds to those patients who do not respond to the drug. Thus there are no optima in this stratum and the corresponding plot in panel B of Figure 4.3 is empty. Panels C and D of Figure 4.3 are not empty for this stratum, however, since $f = 0$ everywhere, and so it is of interest to see how the

algorithms estimate this value. Under this scenario, the personalized algorithm converges to the locations of $\mathbf{d}_{opt,k}$, coming within 1-1.5 dosing units of the optimal depending on the standardized effect size, whereas the standard algorithm does not (Panel B of Figure 4.3). Since the standard algorithm targets the global optimum, it performs best in stratum $(z_1 = 1, z_2 = 1)$ where the standardized effect is the largest. Accurate estimation of f is challenging under this scenario (Panels C and D of Figure 4.3). The personalized algorithm shows evidence of convergence toward the true values of f (i.e., RPSEL and absolute deviation of the posterior mean estimates decreasing to 0) whereas the standard algorithm does not. The standard algorithm yields only a single estimate of f and so must split the difference among the different efficacy function values across the strata. This scenario suggests that as the number of strata grow, and thus also the likelihood for some degree of response heterogeneity to be present, the performance of the standard algorithm will be further degraded. This will depend on the degree of response heterogeneity present, however, and we would expect the standard algorithm to have reasonable performance under scenarios with less extreme response heterogeneity.

In summary, when heterogeneity exists across the strata, the personalized algorithm is superior in both identifying the locations of the $\mathbf{d}_{opt,k}$ and estimating f. When no heterogeneity exists, the standard algorithm is slightly more efficient. Additionally, the proposed methods have performed well for both monotonic and non-monotonic dose-efficacy surfaces, and have done so without utilizing strong prior information.

4.4 Dose-finding design for an intraocular implant

In this section we focus on the intraocular implant example. The goal is to develop an intraocular implant with an optimal dose combination of two agents which reduce intraocular pressure (IOP), a laboratory confirmed measurement. The normal range of IOP is 12-21 mmHg, with 21-30 mmHg considered elevated IOP. Elevated IOP is a risk factor for ocular hypertension and glaucoma, and is strongly associated with increased vision loss (Leske et al.,

2003). The implant seeks to reduce elevated IOP by combining two agents, with doses d_1 and d_2, each of which has been in use individually in topical form for many years. The agents are well tolerated and we do not expect any drug-related adverse events. We are interested in obtaining the optimal dose combination of these two agents using reduction in IOP from baseline as a continuous efficacy measure. It is hypothesized that higher doses do not necessarily imply greater efficacy, which is expected to plateau or even decrease at higher levels of agent concentration. Additionally, we expect response heterogeneity to exist with respect to a particular characteristic of the lens of the eye, which we treat as a binary covariate Z_1, and are interested in a design which permits identification of potentially different optimal dose combinations according to this patient characteristic.

We allow the dose-finding algorithm to explore the dose-combination region, assumed to be standardized, subject to the manufacturing precision constraint of 0.25 standardized dosing units and deem it ethically permissible to proceed without safety-related dose escalation/de-escalation rules. It is hypothesized that the implant can reduce elevated IOP by 5 mmHg in individuals with $Z_1 = 0$, but may be even more effective in individuals with $Z_1 = 1$, leading to reductions as high as 10 mmHg. To assess the cost and size of a hypothetical trial, we are interested in comparing the standard and personalized dosing approaches under different stopping rule specifications for the scenario described above. The final design is then selected as the one which balances good performance while controlling expected cost. Costs are measured in terms of enrolled participants and also the number of unique dose combinations, since there are engineering costs associated with production of novel doses.

The goal is to minimize the efficacy function. The data generating mechanism is $y = f(\mathbf{d}, z) + \epsilon$ where $\epsilon \sim N(0, \sigma_y^2)$, with the specification of $f(\mathbf{d}, z)$ included in the second panel of Table 4.1 (row labeled "Implant") and plotted in panel A of Figure 4.4. We use the same indexing for $\mathbf{d}_{opt,k}$ and $f_{opt,k}$ as described previously. The value of $\sigma_y = 5$ ensures medium standardized

effect sizes of 1 and 2 for $Z_1 = 0$ and $Z_1 = 1$, respectively, across non-monotonic dose-efficacy surfaces.

Two settings of the algorithms are compared for a maximum sample size of 80: one setting includes a higher number of replications at a smaller number of doses, and the other includes a smaller number of replications at a larger number of doses. We denote the standard-/personalized algorithms under the first setting as $S1/P1$ and under the second setting as $S2/P2$. Under the first setting, $S1/P1$ are run under the same specifications described in the previous section, where $r = 4$ and $r = 2$ for the standard versus personalized algorithms, respectively. Under the second setting, $S2/P2$ are run with $r = 2$ and $r = 1$ for the standard versus personalized algorithms, respectively, for $c = 10$ initial dose combinations selected via Sobol sequences.

As the sponsor is concerned about cost and size of the trial, early stopping is permitted using the rule defined in (4.4). Early stopping is investigated by choosing values of δ as previously described such that there is a moderately high chance of stopping after roughly 40 or 60 total participants are enrolled in the trial (denoted by the values of n_{stop} in Figure 4.4). These values are $\{0.00179, 0.000971\}$ for $S1$, $\{0.00670, 0.00345\}$ for $P1$, $\{0.00140, 0.000820\}$ for $S2$, and $\{0.00565, 0.00298\}$ for $P2$. Since we are considering a dual-agent dose combination, $J = 2$ and we thus require the stopping criteria in (4.4) to be satisfied $J + 1 = 3$ times before stopping early.

For the personalized algorithm, we permit stratum specific early stopping. Importantly, should exploration of one stratum stop early, we allocate the remaining budget to recruitment of participants in the other stratum. This assumes the sponsor can target recruitment specifically for this group. This reallocation enables resources to be utilized in strata which are harder to optimize, and so may increase performance. We compare this to no early stopping ($i.e., n_{stop} = 80$), where the numbers of participants enrolled in each stratum are equal. When combined with the two settings for each algorithm, 12 unique designs are defined:

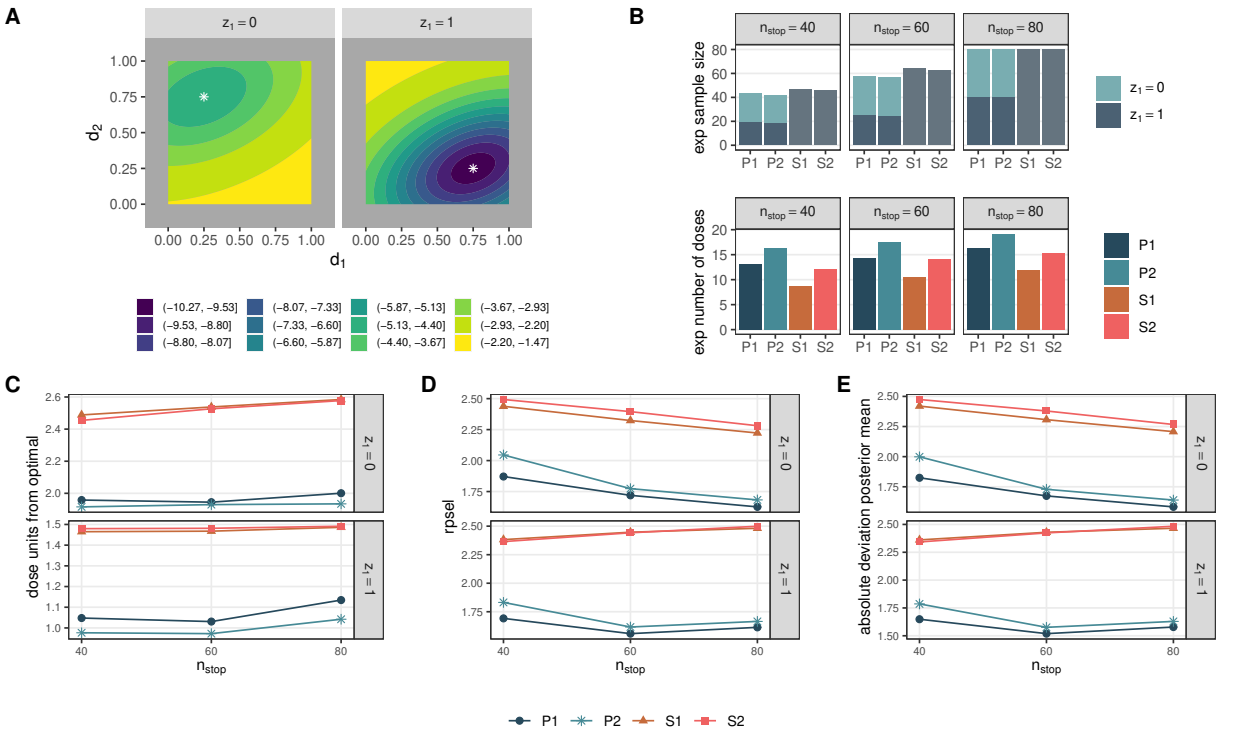

Figure 4.4: Intraocular Implant Scenario: A) Efficacy function, white stars denote $\mathbf{d}_{opt,k}$, B) expected sample size and expected number of unique doses evaluated, C) expected dose units from the optimal dose combination as defined in (4.6) by iteration, D) average RPSEL as defined in (4.7) by iteration, and E) average absolute deviation of the posterior mean estimate by iteration. On the x-axis in (B), P stands for personalized and S for standard, followed by the number referring to high replication for a smaller number of doses (1) or low replication for a larger number of doses (2).

$P1/P2/S1/S2$ each of which has three stopping rules, denoted by $n_{stop} = \{40, 60, 80\}$. All computing and inference is performed as previously described. The performance of the designs is compared using the previously defined criteria which are estimated via 1,000 Monte Carlo replicates.

The expected sample size and expected number of unique doses evaluated by each design for the scenario described above is included in panel B of Figure 4.4. The standard and personalized designs have approximately equal expected samples sizes within each stopping rule, but the personalized designs are expected to evaluate more unique doses on average. The performance of the personalized designs with respect to expected dosing units from the $\mathbf{d}_{opt,k}$ (panel C of Figure 4.4) is comparable within each stopping rule, having differences of

around 0.1 standardized dosing units, which are too small to be practically meaningful.

There is some suggestion of a possible increase in the expected number of dosing units from $\mathbf{d}_{opt,k}$ under no early stopping (i.e., $n_{stop} = 80$) as compared to early stopping. To investigate this, an additional simulation was performed and compared designs with maximum samples sizes of 40 and 60 under no early stopping to those which permit early stopping at roughly these same number of participants. The simulation suggests that the increase mentioned above results from the equal allocation of participants across strata under no early stopping. When early stopping is permitted, a larger proportion of evaluations is allocated to the stratum which is more difficult to optimize, and so provides an improved model fit at each iteration of the algorithm which improves the dose-finding overall. This difference in proportions can be observed in panel B of Figure 4.4 by noting that higher proportions of the expected sample sizes under early stopping rules come from stratum $Z_1 = 0$, which has a smaller standardized effect size and is thus harder to optimize. Regardless, under the current scenario, the observed difference in expected dosing units from $\mathbf{d}_{opt,k}$ between the stopping rules $n_{stop} = 60$ and $n_{stop} = 80$ for the personalized designs is too small to be meaningful. However, future work should more fully investigate how equal versus unequal allocation of participants across the strata at each algorithm iteration impacts design performance.

Design $P1$ estimates f the best, supporting findings in the literature that suggest higher degrees of replication can beneficial for estimation under noisy settings (Binois et al., 2018). This difference is most apparent between designs $P1$ and $P2$ under stopping rule $n_{stop} = 40$ (panels D-E of Figure 4.4). The standard algorithms perform poorly across all performance metrics, recommending doses which are on average farther away from $\mathbf{d}_{opt,k}$, and poorly estimating the efficacy function f at the recommended dose combinations (panels C-E of Figure 4.4). This poor performance is expected since response heterogeneity is present in the true data generating mechanism. If response heterogeneity were not present, we would expect the standard algorithms to be slightly more efficient than their personalized counterparts as

was observed in the simulation study from the last section.

To suggest a final design to the sponsor, we use Figure 4.4 as a visual aid. Since response heterogeneity is expected a priori, the poor performance of the standard designs under this scenario renders them inappropriate. Instead, we select the personalized $P1$ design with $n_{\text{STOP}} = 60$. For roughly the same expected sample size but for fewer unique dose evaluations, this design yields final dose suggestions which are as close to $\mathbf{d}_{opt,k}$ as design $P2$. This design also offers a mild improvement in the estimation of f as compared to $P1$ with $n_{stop} = 40$. Choosing $P1$ with $n_{stop} = 60$ over $P1$ with $n_{stop} = 40$ does come with additional cost, however: the design with $n_{stop} = 40$ expects to evaluate 13 unique doses and enroll approximately 44 participants, whereas the design with $n_{stop} = 60$ expects to evaluate about 15 unique dose combinations (a 15% increase) and enroll approximately 58 participants (a 32% increase). The sponsor would need to weigh the increased engineering and enrollment costs against the increase in performance.

4.5 Discussion

In this manuscript, we proposed the use of Bayesian optimization for early phase multi-agent dose-finding trials in a tolerated toxicity setting. We showed the benefit of taking a personalized approach for dual-agent trials when heterogeneity exists across a set of prespecified subgroups. As expected, under no response heterogeneity the personalized approach is slightly less efficient. As noted in the introduction, parametric models may suffer from the curse of dimensionality when transitioning from standard to personalized dose-finding, as they require terms for all higher order dose-covariate interactions. By using the anisotropic squared exponential GP kernel in the Bayesian optimization methods proposed here, however, only a single additional parameter per covariate is required (the additional lengthscale parameter corresponding to that covariate). Thus, the proposed methods highlight the benefit and feasibility of adopting a personalized approach toward early phase multi-agent dose-finding trials for both monotonic and non-monotonic dose-response surfaces.

The proposed approach is not without limitations. First, the methods proposed in this work assumed no meaningful toxicity across the dose combination space. Extension to higher-grade toxicity settings through incorporation of dose escalation/de-escalation remains as future work. Second, the personalized approach was demonstrated by considering dual-agent dose combinations in predefined subgroups only. Extension to dose combinations with more than two agents is trivial, but extension to continuous covariates without categorization is not, and could be the subject of future investigations. Another direction for future development is to extend the proposed approach to binary and ordinal outcomes where the proposed response models may be defined over a latent continuous surface.

In this manuscript, Bayesian optimization is utilized as a global optimizer. While other global optimization methods exist (e.g., genetic algorithms and simulated annealing), they require many function evaluations and are thus not appropriate for early phase dose-finding trials where evaluations are expensive (Bull, 2011; Tracey and Wolpert, 2018). We employed the AEI acquisition function under a GP surrogate model. Performance of the algorithms under additional acquisition functions, surrogate models, and/or kernel functions should be investigated. The (stationary) separable anisotropic kernel used in this work assumes that all strata have the same correlation structure and that the covariance between points in different strata are changed by a multiplicative factor only, which may not be true in general. Indeed, under simulation scenario 3 which included two binary covariates, the efficacy function is zero everywhere for stratum $(z_1 = 0, z_2 = 0)$, so this assumption is not true in this case. The efficacy function values in this stratum are perfectly correlated, whereas those corresponding to dose combinations in other strata are not. Future work should investigate relaxing the assumption of stationarity by using kernels that are non-separable (e.g., including dose-covariate interaction terms in the kernel function (4.5)), or even non-stationary, or deep GPs (Sauer et al., 2023). Finally, if the number of included covariates is large and there is reason to believe that only low level interactions between the drug combinations and covariates exist, different surrogate models could be employed, such as additive GPs (Duvenaud et al.,

2011) or Bayesian additive regression trees (Chipman et al., 2010).

The simulations performed in this manuscript assumed that the subgroups represented by each strata were equally prevalent and had the same enrollment rates throughout the trial. This represents an idealized situation and may not be reasonable to assume in practice. In the extreme scenario of one strata having zero patients, the algorithm would simply base recommendations for this stratum on prior information only. Furthermore, due to the borrowing of information across the strata, the estimation for sparse strata may be dominated by strata with many patients. For these reasons, Zhang et al. (2024) suggest subgroup-specific dose-optimization only be performed for pre-defined subgroups which have large enough sample sizes in the trial, a recommendation which supports the simulation scenarios evaluated in the present work.

Finally, we adopted an empirical Bayes approach toward the GP hyperparameter estimation to decrease the computational burden of the simulations. Likelihood methods can yield poor results when the sample sizes are small, as in early phase dose-finding trials, and so full Bayesian inference may be preferred (Bull, 2011; Wang and de Freitas, 2014). Unfortunately, the Markov chain Monte Carlo methods typically used to perform full Bayesian inference are prohibitively expensive for the algorithms proposed here, and so a sequential Monte Carlo approach may be a less computationally demanding alternative (Gramacy and Polson, 2011).

Data availability

No new data were created or analyzed in this study. The R scripts used for the simulations and graphics can be found on a public GitHub repository at `https://github.com/jjwillard/bayesopt_pers_df`.

Acknowledgements

JW acknowledges the support of a doctoral training scholarship from the Fonds de recherche du Québec - Nature et technologies (FRQNT) and a research exchange internship with PharmaLex Belgium. SG acknowledges the support by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada (NSERC) and a Chercheurs-boursiers (Junior 1) award from the Fonds de recherche du Québec - Santé (FRQS). EEMM acknowledges support from a Discovery Grant from NSERC. EEMM is a Canada Research Chair (Tier 1) in Statistical Methods for Precision Medicine and acknowledges the support of a chercheur de mérite career award from the FRQS. This research was enabled in part by support provided by Calcul Québec and the Digital Research Alliance of Canada.

Chapter 5

Bayesian optimization for identification of optimal biological dose combinations in personalized dose-finding trials

Preamble to Manuscript 3. In this final manuscript, I generalize the Bayesian optimization-based approach to personalized dose-finding for combination therapies proposed in the previous chapter to the setting of higher-grade toxicities. The goal is now to find patient-specific optimal biological dose combinations, defined as safe dose combinations which maximize therapeutic benefit for an individual patient. Under this setting, toxicity information must be considered in the dose-finding procedure, which is viewed as a constrained optimization problem. Separate Gaussian processes are used to model the efficacy and toxicity functions under the assumption of their conditional independence given dose. To collect the initial responses while minimizing the number of toxic doses that are administered, a dose-escalation scheme is proposed. A sequential search policy is then defined by optimizing a constrained version of a commonly used acquisition function. The performance of the procedure is assessed via a simulation study under multiple scenarios.

This work is motivated by a trial design for a combination therapy used to treat obstructive sleep apnea, a highly prevelant disease for which no pharmacotherapy exists (Benjafield et al., 2019). Response heterogeneity is expected to exist with respect to sleep apnea severity and, to allow for a greater degree of benefit, a higher level of toxicity should be permitted for the most severe subtype. The methodological contributions of this manuscript are i) to generalize the personalized dose-finding approach for combination therapies proposed in Chapter 4 to the setting of higher grade toxicities, ii) to accommodate patient-specific heterogeneity with respect to both the efficacy and toxicity functions and to the permissible levels of toxicity, and iii) to propose a dose-escalation scheme which reduces the number of toxic doses administered to the initial patients. The substantive contribution of this work is to demonstrate the feasibility of personalized dose-finding designs for combination therapies under higher-grade toxicity settings despite the limited sample sizes in early phase trials.

The corresponding manuscript has been submitted for review to a statistical journal.

Bayesian Optimization for Identification of Optimal Biological Dose Combinations in Personalized Dose-Finding Trials

James Willard[1], Shirin Golchi[1], and Erica E.M. Moodie[1]

[1]Department of Epidemiology, Biostatistics, and Occupational Health, McGill University

Abstract

Early phase, personalized dose-finding trials for combination therapies seek to identify patient-specific optimal biological dose (OBD) combinations, which are defined as safe dose combinations which maximize therapeutic benefit for a specific covariate pattern. Given the small sample sizes which are typical of these trials, it is challenging for traditional parametric approaches to identify OBD combinations across multiple dosing agents and covariate patterns. To address these challenges, we propose a Bayesian optimization approach to dose-finding which formally incorporates toxicity information into both the initial data collection process and the sequential search strategy. Independent Gaussian processes are used to model the efficacy and toxicity surfaces, and an acquisition function is utilized to define the dose-finding strategy and an early stopping rule. This work is motivated by a personalized dose-finding trial which considers a dual-agent therapy for obstructive sleep apnea, where OBD combinations are tailored to obstructive sleep apnea severity. To compare the performance of the personalized approach to a standard approach where covariate information is ignored, a simulation study is performed. We conclude that personalized dose-finding is essential in the presence of heterogeneity.

5.1 Introduction

Early phase clinical trials assess the safety and efficacy of first-in human doses of experimental therapies. Bayesian adaptive designs, which utilize posterior or posterior predictive distributions for sequential decision making, are commonly employed for these trials, and can be divided into three major families: algorithmic, model-assisted, and model-based (Berry et al., 2010). Here, our focus is on model-based designs, which posit a full probability model for the dose-response surface of interest. Additionally, they permit clinicians to formally encode prior beliefs in the trial, which may be beneficial given the small sample sizes characteristic of early phase trials. Much of the historical work in early phase designs is motivated by the development of cytotoxic agents in oncology, where the goal is to find the maximum tolerated dose (MTD) while typically assuming monotonic dose-response surfaces and binary endpoints (O'Quigley et al., 1990; Babb et al., 1998).

Recently, interest has grown in molecularly targeted compounds which may display benefit at dosages lower than those that induce toxicity or the MTD (Le Tourneau et al., 2009). In these trials, traditional toxicity responses have been replaced by other drug-related biological effects, which are often measured on a continuous scale, e.g., plasma drug concentration and measures of target inhibition in tissues of interest (Korn, 2004; Le Tourneau et al., 2009). Additionally, many binary toxicity responses arise from dichotomizing a continuous measure, which is well known to lead to loss of information, decreased efficiency, increased sample sizes, and may introduce bias (Farewell et al., 2004; Royston et al., 2006; Wason et al., 2011). These developments have increased interest in continuous toxicity responses (e.g., Chen et al. (2010); Lee et al. (2012); Le-Rademacher et al. (2020)). Despite this, the literature for dose-finding using continuous responses remains less developed than that using binary responses. With the development of novel therapies where it is possible that the MTD is not the most efficacious dose, interest has grown in designs which incorporate both efficacy and toxicity information to find an optimal biological dose (OBD; e.g., Thall and

Cook (2004); Yuan et al. (2016b)). Furthermore, since monotonicity of the dose-response surfaces depends on the type of therapy being investigated and may not hold in general (Li et al., 2017), flexible designs have been proposed to identify OBD combinations under non-monotonicity (e.g., Houede et al. (2010); Mozgunov and Jaki (2019b)) and for continuous dose-response surfaces (Mozgunov and Jaki, 2019a).

All dose-finding designs mentioned above assume the same dose is optimal for every patient in the population. We refer to this "one-size-fits-all" approach as *standard dose-finding*. With the continued identification and development of novel biomarkers, interest has grown in more personalized approaches to medicine. *Personalized dose-finding* seeks to find optimal doses based on individual patient characteristics. Despite this growing interest, the literature for personalized dose-finding trials remains underdeveloped as the limited sample sizes in these trials make it challenging to extend many parametric dose-finding methods to the personalized setting, where potentially many dose-covariate interaction terms must be estimated. Personalized dose-finding trials for monotherapy have been explored previously (e.g., Babb and Rogatko (2001); Guo and Zang (2022)). Our focus in this manuscript is on personalized dose-finding trials for combination therapies, where the additional dosing agents increase the dimensionality and exacerbate the estimation challenges. Mozgunov et al. (2022) investigated dual-agent personalized dose-finding with an application to an opiate detoxification trial, where the level of one of the agents was pre-specified externally by clinicians. Willard et al. (2023) proposed a flexible personalized dose-finding design for continuous responses which uses Bayesian optimization to explore all dosing dimensions under no monotonicity assumptions. However, they assumed a minimal toxicity setting where every dose combination was assumed to be safe, and so formal dose escalation rules were not considered. The present work extends this approach to the setting of higher-grade continuous toxicities by proposing a dose escalation scheme to collect the initial data, and by formally incorporating the toxicities into the search strategy employed by the Bayesian optimization based dose-finding method.

This work is motivated by a dose-finding design for a combination therapy which treats obstructive sleep apnea (OSA). OSA is a common sleep disorder which is estimated to affect over 930 million adults worldwide (Benjafield et al., 2019). Continuous positive airway pressure (CPAP) is the standard treatment since it is very effective, but it is often poorly tolerated and results in low patient adherence (Rotenberg et al., 2016). There is currently no approved pharmacotherapy for OSA, though several recent studies have investigated the combination of antimuscarinic agents (oxybutynin and aroxybutynin) and a norepinephrine reuptake inhibitor (atomoxetine) as therapy (Schweitzer et al., 2023; Aishah et al., 2023; Rosenberg et al., 2022). While these studies suggest that combination therapy is effective in reducing a key continuous endpoint, they revealed that potential response heterogeneity exists with respect to OSA severity and so a more targeted therapeutic approach would be beneficial. The current work aims to develop a design that can identify OBD combinations which are tailored to OSA severity.

The remainder of this manuscript is organized in the following manner. We first provide a brief introduction to Bayesian optimization and review the dose-finding approaches established for continuous responses in Willard et al. (2023). Next, we propose a formal dose-escalation scheme for continuous toxicities, which generalizes these approaches to the setting of higher-grade toxicities. This is followed by a simulation study where we compare the performance of the proposed approaches under two scenarios. We then consider the previously described dose-finding design for a dual-agent therapy to treat OSA. We conclude with a discussion.

5.2 Bayesian optimization for dose-finding

Personalized dose-finding Consider a set of P discrete covariates $\mathbf{Z} = \{Z_p\}_{p=1}^{P}$, and define K strata as the Cartesian product of their levels. Personalized dose finding seeks to identify optimal dose combinations for each of the K strata. We refer to the continuous therapeutic function of interest (i.e., efficacy or utility) as the *efficacy* function, denoted by $f(\mathbf{d}, \mathbf{z}) \in \mathbb{R}$,

and define the continuous toxicity function as $g(\mathbf{d}, \mathbf{z}) \in \mathbb{R}$, where both are assumed to have been transformed such that smaller values denote being more desirable. Our goal is to find, for each stratum k, the combination of $J \in \mathbb{N}$ dosing agents $\mathbf{d} = (d_1, ..., d_J) \in \mathbb{D} \subset \mathbb{R}^J$ which minimizes $f(\mathbf{d}, \mathbf{z}_k)$, subject to a tolerable level of toxicity $g_k^\dagger$:

$$\underset{\mathbf{d} \in \mathbb{D}}{\operatorname{argmin}} f(\mathbf{d}, \mathbf{z}_k) \quad \text{subject to} \quad g(\mathbf{d}, \mathbf{z}_k) \leq g_k^\dagger \text{ for } \quad k = 1, ..., K.$$

Unique toxicity thresholds $g_k^\dagger$ allow permissible toxicity to be defined with respect to specific covariate profiles. This is useful in settings where higher levels of toxicity may be permitted for more severe disease subtypes, for example. We solve the above minimization problem using Bayesian optimization, a derivative-free method which finds the global optima of expensive-to-evaluate objective functions (Garnett, 2023). Bayesian optimization relies on stochastic surrogate models, commonly Gaussian Processes (GP), to estimate the efficacy and toxicity functions, and then employs an acquisition function to define a sequential search policy. It has been historically used for global optimization problems within the engineering, machine learning, and computer experiments literature (e.g., Kushner (1964); Zhilinskas (1975); Jones et al. (1998); Gramacy (2020)), and has seen recent application in early phase dose-finding trials (Takahashi and Suzuki, 2021b,0; Willard et al., 2023). Below, we generalize the approach proposed in Willard et al. (2023) to higher-grade toxicity settings, by using Bayesian optimization for dose-finding in the presence of toxicity constraints.

To start, $n = \sum_{k=1}^{K} r_k \times c_k$ initial responses are collected, where for the k^{th} stratum, r_k independent patients are evaluated at c_k initial dose combinations. A proposed method for collecting these initial responses under uncertain toxicity is described in detail further below. These initial responses yield noisy observations of both the efficacy function, $y_f = f(\mathbf{d}, \mathbf{z}) + \epsilon_f$, and the toxicity function, $y_g = g(\mathbf{d}, \mathbf{z}) + \epsilon_g$, where $\epsilon_f \sim N(0, \sigma_{y_f}^2)$, $\epsilon_g \sim N(0, \sigma_{y_g}^2)$. We denote the vector of responses for the patient i by $\mathbf{y}_i = (y_f, y_g)_i^T$, the vector of efficacy responses for all patients by $\mathbf{y}_f = (y_{f,1}, ..., y_{f,n})^T$, and the vector of toxicity responses for all patients by $\mathbf{y}_g = (y_{g,1}, ..., y_{g,n})^T$. In many contexts, it is reasonable to assume conditional independence

between the efficacy and toxicity functions given dose, and so we do throughout. We describe how to relax this assumption in the discussion section for scenarios where they are expected to remain correlated after adjusting for dose. Independent GP priors are placed on the efficacy and toxicity functions, the details of which are described below for the efficacy function only since those for toxicity are similar. In cases where it is reasonable to assume the dose-toxicity function is monotonically non-decreasing, GP models incorporating monotonicity constraints (Lin and Dunson, 2014; Golchi et al., 2015) or less flexible models could be utilized, though we do not consider these approaches here. The dose-toxicity function may be non-monotonic in cases where a toxicity score is defined as an aggregate of adverse events (AEs; e.g., Le-Rademacher et al. (2020)), where different dose combinations yield different distributions of the individual AEs that, when aggregated, potentially yield a non-monotonic dose-toxicity surface. We utilize the following unconstrained GP model for the efficacy function:

$$f(\mathbf{d}, \mathbf{z}) \sim GP\left(m_f(\mathbf{d}, \mathbf{z}), \nu_f \mathcal{K}_f((\mathbf{d}, \mathbf{z}), (\mathbf{d}', \mathbf{z}'))\right)$$

where $m_f(\mathbf{d}, \mathbf{z})$ is the mean function, and $\mathcal{K}_f((\mathbf{d}, \mathbf{z}), (\mathbf{d}', \mathbf{z}'))$ is a correlation function (kernel) multiplied by scale parameter ν_f (Binois and Gramacy, 2021). The scale parameter determines the variability of the efficacy function throughout the dose combination space. We utilize GP models with a constant mean function, $m_f(\mathbf{d}, \mathbf{z}) = \beta_f$, though note that additional information about the efficacy function (e.g., pharmacokinetic/pharmacodynamic models) or toxicity function (e.g., quantitative systems toxicology models) can be incorporated into the mean function. We employ a stationary anisotropic squared exponential kernel,

$$\mathcal{K}_f((\mathbf{d}, \mathbf{z}), (\mathbf{d}', \mathbf{z}')) = \exp\left\{-\left(\sum_{j=1}^{J} \frac{(d_j - d_j')^2}{2l_{f,j}^2} + \sum_{p=1}^{P} \frac{(z_p - z_p')^2}{2l_{f,p}^2}\right)\right\} \tag{5.1}$$

parameterized by characteristic length-scales $\{l_{f,1}, ..., l_{f,P}\}$. These control how quickly the correlation between two dose combinations decays with respect to each dosing agent and across covariate patterns. Throughout this work, we assume stationary dose-response surfaces, where the degree of correlation between two dose combinations depends only on their

distance from one another in the input space. We comment on relaxing this assumption in the discussion.

Specifying this GP prior induces a multivariate normal distribution on the efficacy function observations, $\mathbf{y}_f \sim N(\beta_f \mathbf{1}_n, \nu_f \mathbf{K}_f)$, where $\mathbf{K}_f(i,j) = \mathcal{K}_f((\mathbf{d}_i, \mathbf{z}_i), (\mathbf{d}_j, \mathbf{z}_j)) + \tau_f^2 \mathbb{1}_{i=j}$ and τ_f^2 is a noise parameter. After data $\mathcal{D} = \{(\mathbf{d}_i, \mathbf{z}_i, \mathbf{y}_i)\}_{i=1}^n$ are observed, the kernel hyperparameters $\boldsymbol{\theta} = \{\nu_f, \tau_f^2, l_{f,1}, ..., l_{f,P}\}$ are estimated via maximum likelihood and replaced by their point estimates, taking an empirical Bayes approach (Binois and Gramacy, 2021).

This yields the posterior distribution of the efficacy function at a new dose combination $\widetilde{\mathbf{d}}$ in stratum k, denoted by $\widetilde{\mathbf{d}}_k = (\widetilde{\mathbf{d}}, \widetilde{\mathbf{z}}_k)$, as $p(f \mid \mathcal{D}, \widetilde{\mathbf{d}}_k) = N(\mu_f(\widetilde{\mathbf{d}}_k), \sigma_f^2(\widetilde{\mathbf{d}}_k))$ (Binois and Gramacy, 2021), such that

$$\begin{aligned} \mu_f(\widetilde{\mathbf{d}}_k) &= \hat{\beta}_f + \mathbf{k}_f(\widetilde{\mathbf{d}}_k)^T \mathbf{K}_f^{-1}(\mathbf{y}_f - \hat{\beta}_f \mathbf{1}) \\ \sigma_f^2(\widetilde{\mathbf{d}}_k) &= \nu_f \mathcal{K}_f(\widetilde{\mathbf{d}}_k, \widetilde{\mathbf{d}}_k) - \nu_f \mathbf{k}_f(\widetilde{\mathbf{d}}_k)^T \mathbf{K}_f^{-1} \mathbf{k}_f(\widetilde{\mathbf{d}}_k) + \frac{(1 - \mathbf{k}_f(\widetilde{\mathbf{d}}_k)^T \mathbf{K}_f^{-1} \mathbf{1})^2}{\mathbf{1}^T \mathbf{K}_f^{-1} \mathbf{1}} \\ \hat{\beta}_f &= \frac{\mathbf{1}^T \mathbf{K}_f^{-1} \mathbf{y}_f}{\mathbf{1}^T \mathbf{K}_f^{-1} \mathbf{1}} \end{aligned} \tag{5.2}$$

where $\mathbf{K}_f$ is $n \times n$, $\mathbf{k}_f(\widetilde{\mathbf{d}}_k) = [\mathcal{K}_f((\mathbf{d}_1, \mathbf{z}_1), \widetilde{\mathbf{d}}_k), ..., \mathcal{K}_f((\mathbf{d}_n, \mathbf{z}_n), \widetilde{\mathbf{d}}_k)]^T$ is $n \times 1$, and $\hat{\beta}_f$ is a plug-in estimate for the mean. The posterior distribution of the toxicity function at $\widetilde{\mathbf{d}}$ within stratum k is obtained similarly and has the same functional form as above but where f is replaced by g for all observations and kernel hyperparameters.

The next dose combination $\mathbf{d}_k^{(c_k+1)}$ within stratum k, is then selected as that which maximizes an acquisition function, denoted by $\alpha(\widetilde{\mathbf{d}}_k \mid \mathcal{D})$:

$$\mathbf{d}_k^{(c_k+1)} = \underset{\widetilde{\mathbf{d}} \in \mathbb{D}}{\operatorname{argmax}}\, \alpha(\widetilde{\mathbf{d}}_k \mid \mathcal{D}).$$

One possible acquisition function which can be used in the presence of toxicity constraints is a constrained version of the Expected Improvement (cEI), defined as

$$\alpha_{cEI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D}) = \mathbb{E}\left[\max(0, f_k^* - f(\widetilde{\mathbf{d}}_k)) \mathbb{1}\{g(\widetilde{\mathbf{d}}_k) \leq g_k^\dagger\} \mid \mathcal{D}, \widetilde{\mathbf{d}}_k\right],$$

and available in closed form under the assumption of conditional independence between the efficacy and toxicity functions given dose (Jones et al., 1998; Gardner et al., 2014):

$$\alpha_{cEI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D}) = \left[(f_k^* - \mu_f(\widetilde{\mathbf{d}}_k))\Phi\left(\frac{f_k^* - \mu_f(\widetilde{\mathbf{d}}_k)}{\sigma_f(\widetilde{\mathbf{d}}_k)}\right) + \sigma_f(\widetilde{\mathbf{d}}_k)\phi\left(\frac{f_k^* - \mu_f(\widetilde{\mathbf{d}}_k)}{\sigma_f(\widetilde{\mathbf{d}}_k)}\right)\right] \Phi\left(\frac{g_k^\dagger - \mu_g(\widetilde{\mathbf{d}}_k)}{\sigma_g(\widetilde{\mathbf{d}}_k)}\right).$$

Above, f_k^* denotes the current optimum of the efficacy function in stratum k, $g_k^\dagger$ denotes the stratum-specific toxicity constraint, $\Phi(\cdot)$ and $\phi(\cdot)$ denote the standard normal cumulative distribution function and probability density function, respectively, and $\mu(\widetilde{\mathbf{d}}_k)$ and $\sigma(\widetilde{\mathbf{d}}_k)$ are given above in (5.2). The expression inside the brackets serves to balance the trade-off between exploring regions of the dose combination space where the efficacy function is imprecisely estimated, and exploiting regions which have desirable values of the efficacy function. This trade-off is weighted by the posterior probability of satisfying the toxicity constraint, giving higher weight to points which are more likely to be safe. Under a noisy setting, we set $f_k^* = \min_{\widetilde{\mathbf{d}} \in \mathbb{A}_k} \boldsymbol{\mu}_f(\widetilde{\mathbf{d}}_k)$, the minimum of the posterior mean of the efficacy function in the set of safe doses, $\mathbb{A}_k$, in stratum k. A dose combination is considered to be safe if $P(g(\widetilde{\mathbf{d}}_k) \leq g_k^\dagger \mid \mathcal{D}) > \gamma_k$, where γ_k is a tuning parameter which can be specified through a sensitivity analysis. Larger values of γ_k imply a more stringent definition of safety whereas smaller values imply a more permissive definition. In the event where $\mathbb{A}_k = \emptyset$ during a single iteration, f_k^* is set as the posterior mean of the maximizer of this probability statement (i.e., the average efficacy at the dose predicted to be safest; Gelbart et al. (2014)). We note that dose-finding in stratum k terminates before assigning $\mathbf{d}_k^{(c_k+1)}$ if no dose combinations are predicted to be safe for several consecutive iterations. Since observations of the toxicity function are noisy, random variation in the data may cause incorrect early stopping if it is permitted after only a single violation of this probability statement. Thus, we require the probability statement to be violated for $J+1$ iterations before termination, where J is the number of dosing agents, following a recommendation for stopping rules under noise (Huang et al., 2006).

After observing responses for r_k new patients at $\mathbf{d}_k^{(c_k+1)}$, the data are updated and the GP

models are refit to obtain new posterior distributions for f and g. Then $s = 1, ..., S$ samples $f^{(s)}(\widetilde{\mathbf{d}}_k)$ and $g^{(s)}(\widetilde{\mathbf{d}}_k)$ from these posteriors are obtained to yield S samples from the posterior of the optimal dose combination $p(\mathbf{d}_{opt,k} \mid \mathcal{D})$ as:

$$\mathbf{d}_{opt,k}^{(s)} = \underset{\widetilde{\mathbf{d}} \in \mathbb{A}_k^{(s)}}{\operatorname{argmin}} f^{(s)}(\widetilde{\mathbf{d}}_k)$$

where the admissible set of safe doses in stratum k is defined as $\mathbb{A}_k^{(s)} = \{\widetilde{\mathbf{d}} : P(g^{(s)}(\widetilde{\mathbf{d}}_k) \leq g_k^\dagger) > \gamma_k\}$. Dose-finding within each stratum continues until the sample size limit is reached or a stratum-specific early stopping rule is satisfied. We permit early stopping in stratum k when $\max_{\widetilde{\mathbf{d}} \in \mathbb{D}} \alpha_{cEI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D}) < \delta_k$. Due to random variability in the data, we require this statement to be satisfied for at least $(J + 1)$ consecutive iterations, where δ_k is calibrated through sensitivity analysis. For example, early stopping would occur at iteration 7 if $\delta_k = 0.11$ and $\max_{\widetilde{\mathbf{d}} \in \mathbb{D}} \alpha_{cEI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D})$ is $0.12, 0.10, 0.09, 0.08$ for iterations $4, 5, 6, 7$, respectively. For convenience below, we refer to this as early stopping for "efficacy", noting that it is technically early stopping due to minimal expected improvement over the current best observation, which need not imply a significant therapeutic response. The sequential procedure described above is referred to as the *personalized optimization algorithm*. Note that in the present work and in contrast to the personalized optimization algorithm in Willard et al. (2023), we explicitly model toxicity and incorporate it into the sequential search policy defined by $\alpha_{cEI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D})$. We note that the *standard optimization algorithm* is a special case of the personalized approach, where covariate information is ignored (i.e., $\mathbf{Z} = \emptyset$) and only the dosing agents themselves are considered for optimization.

5.3 Initial data collection

Expansion of dose combination region It is important to limit to the extent possible the number of patients who experience adverse events and who are assigned to toxic doses. After the initial data are collected and the dose-finding algorithm described above begins, the achievement of these objectives is controlled through the cEI acquisition function, which gives

lower values to doses which are predicted to be toxic, making them less likely to be evaluated. To limit toxicities among the initial patients, however, doses must be explored in a sequential manner, where lower (i.e., less toxic) doses are evaluated before higher (i.e., more toxic) doses. In a monotherapy setting, this is straightforward since there is a natural ordering from lower to higher doses. In a combination therapy setting, this is more challenging since many dose combinations may yield comparable toxicity and so there is no longer a natural ordering. One challenge in designing a dose-escalation scheme under this setting is in deciding which doses should be candidates for evaluation at each stage of the escalation. For dual-agent combination therapies, Sweeting and Mander (2012) showed that the MTD was more efficiently identified when, at each iteration of the dose escalation algorithm, the levels of all dosing agents were permitted to be increased rather than the level of a single agent only. This amounts to allowing dose escalation to proceed along the "diagonal" when considering a two-dimensional dosing grid. We adopt a similar idea below by permitting our continuous dose combination region of interest in stratum k, denoted by $\mathbb{D}_k$, to be expanded in all coordinate directions at each iteration of dose escalation. We note that expansion of $\mathbb{D}_k$ does not control which doses within this region will be evaluated. That choice is controlled separately using the cEI acquisition function as previously described. However, $\mathbb{D}_k$ serves to restrict the choice of doses the cEI may consider at each iteration, thereby ensuring the initial doses are explored in an escalating fashion.

Consider a continuous dose combination region within stratum k, which we assume to be standardized and bounded, $\mathbb{D}_k \subseteq [0,1]^J$. At the beginning of the dose escalation scheme, we define $\mathbb{D}_k$ to be the lowest standardized dose only, denoted by $\mathbf{d}_k^{(0)} = \mathbf{0}$. We start by evaluating this dose, after which $\mathbb{D}_k$ is expanded. At each iteration of the dose escalation scheme, the sequentially expanded $\mathbb{D}_k$ takes the form of a lower half-space intersected with $[0,1]^J$. To define this intersection, we select a value ρ_k which represents the maximum amount by which a single coordinate dosing dimension may be increased. Then at iteration q, we set $\mathbb{D}_k^{(q)} = \{\mathbf{d} \in [0,1]^J : \mathbf{w}_k^T \mathbf{d} \leq \rho_k \times q, \mathbf{w}_k = \mathbf{1}_J\}$. Smaller values of ρ_k imply a

slower expansion of $\mathbb{D}_k$ whereas larger values imply a faster expansion of $\mathbb{D}_k$. This rate of expansion should depend on the potential severity of toxicities considered, with slower expansion being necessary for contexts which are expected to contain more severe toxicities. We choose $\mathbf{w}_k$ to be the J-dimensional vector of ones to perform uniform expansion across all dosing dimensions. This could be relaxed if prior information suggests non-uniform expansion should be permitted, e.g., if one dosing agent is known to be much safer than the others. We note that expansion of $\mathbb{D}_k$ will continue to eventually include the entire unit hypercube, even if there are regions which are predicted to be overly toxic. Expansion is permitted because exploration within the space is controlled by the cEI acquisition function, which avoids toxic regions and will terminate the dose-finding algorithm in the event of there being no safe doses. As mentioned above, the expansion of $\mathbb{D}_k$ serves only to restrict which doses the cEI can explore at each iteration, ensuring doses are evaluated in an escalating manner.

As an example of this sequential expansion of $\mathbb{D}_k$, consider a dual-agent dose combination setting with standardized doses in $[0,1]^2$ and where $\mathbf{d}^{(0)} = (0,0)$ and $\rho_k = 0.25$. After evaluating $\mathbf{d}_k^{(0)}$, we consider the first iteration of expansion and set $q = 1$. Then we define $\mathbb{D}_k^{(1)} = \{\mathbf{d} \in [0,1]^J : \mathbf{w}_k^T\mathbf{d} \leq 0.25, \mathbf{w}_k = \mathbf{1}_J\}$. That is, at the first iteration, $\mathbb{D}_k$ becomes the triangle defined by vertices $\{(0,0),(0,0.25),(0.25,0)\}$ from which the cEI determines $\mathbf{d}^{(1)}$. At the second iteration of expansion, we set $q = 2$ and define $\mathbb{D}_k^{(2)} = \{\mathbf{d} \in [0,1]^J : \mathbf{w}_k^T\mathbf{d} \leq 0.5, \mathbf{w}_k = \mathbf{1}_J\}$. Thus, at the second iteration, $\mathbb{D}_k$ becomes the triangle defined by vertices $\{(0,0),(0,0.5),(0.5,0)\}$ from which the cEI determines $\mathbf{d}^{(2)}$. This expansion continues until $\mathbb{D}_k$ is defined as the entire unit hypercube, or until the dose-finding algorithm terminates.

To further promote exploration of the dose combination region during the expansion of $\mathbb{D}_k$, regions which are close to previously evaluated doses may be removed. That is, until $\mathbb{D}_k = [0,1]^J$, a modified region can be defined at each iteration as $\mathbb{D}_k^{(q)*} = \mathbb{D}_k^{(q)} \setminus \mathbb{C}_k^{(q)}$, where $\mathbb{C}_k^{(q)} = \bigcup_{i=0}^{q-1} \mathcal{N}_k^{(i)}$ is the set of regions or neighbourhoods $\mathcal{N}_k^{(i)}$ (e.g., hypercubes defined

Table 5.1: Simulation scenarios considered. The data generating mechanism for each scenario is $y_f = f(\mathbf{d}, z_1)+\epsilon_f$ and $y_g = g(\mathbf{d}, z_1)+\epsilon_g$ where $\epsilon_f \sim N(0, \sigma^2_{y_f})$ and $\epsilon_g \sim N(0, \sigma^2_{y_g})$. The table columns contain the location of the optimal dose combination ($\mathbf{d}_{opt}$), the value of the efficacy and toxicity functions at $\mathbf{d}_{opt}$ (f_{opt}/g_{opt}), and the standardized effect sizes (ses_f/ses_g).

	$f(\mathbf{d}, z_1)$	$g(\mathbf{d}, z_1)$	σ_{y_f}	σ_{y_g}	z_1	$\mathbf{d}_{opt}$	f_{opt}	g_{opt}	ses_f	ses_g
Simulation Study	1) $-\mathbb{1}\{z_1 = 0\} \times h(\boldsymbol{\mu}_1) -$ $\mathbb{1}\{z_1 = 1\} \times h(\boldsymbol{\mu}_1)$	$\mathbb{1}\{z_1 = 0\}\times h(\boldsymbol{\mu}_2)+$ $\mathbb{1}\{z_1 = 1\} \times h(\boldsymbol{\mu}_2)$	1.59	0.13	0	(0.5, 0.5)	-1.59	0.13	1	1
					1	(0.5, 0.5)	-1.59	0.13	1	1
	2) $-\mathbb{1}\{z_1 = 0\} \times h(\boldsymbol{\mu}_3) -$ $\mathbb{1}\{z_1 = 1\} \times h(\boldsymbol{\mu}_4)$	$\mathbb{1}\{z_1 = 0\}\times h(\boldsymbol{\mu}_5)+$ $\mathbb{1}\{z_1 = 1\} \times h(\boldsymbol{\mu}_6)$	1.59	0.13	0	(0.25, 0.75)	-1.59	0.13	1	1
					1	(0.75, 0.25)	-1.59	0.13	1	1
OSA	1) $\beta_{z_1,0}+\beta_{z_1,1}d_1+\beta_{z_1,2}d_2+$ $\beta_{z_1,3}d_1d_2+\beta_{z_1,4}d_1^2+$ $\beta_{z_1,5}d_2^2+\beta_{z_1,6}d_1^2d_2^2$	$\theta_0+\theta_1d_1+\theta_2d_2+$ $\theta_3d_1d_2+\theta_4d_1^2+$ $\theta_5d_2^2+\theta_6d_1^2d_2^2$	7.68	1.29	0	(2.5, 75)	-7.68	1.29	1	1
					1	(5, 75)	-13.20	1.63	1.72	1.26

by a small side length l_k, where $l_k \ll \rho_k$) which surround the previously evaluated dose combinations $\{\mathbf{d}_k^{(i)}\}_{i=0}^{q-1}$. Using $\mathbb{D}_k^{(q)*}$ instead of $\mathbb{D}_k^{(q)}$ promotes exploration of the input space during the initial stages of dose-finding, and may prevent the algorithm from getting stuck in certain regions. This can lead to improved GP model fits at early iterations which may improve algorithm performance overall.

5.4 Simulation study

In this section, we perform a simulation study to compare the performance of the personalized and standard approaches for a dual-agent (i.e., $J = 2$) therapy under two scenarios which consider a single binary covariate Z_1. The first scenario considers no response heterogeneity across the strata with respect to the objective and toxicity surfaces, whereas the second considers response heterogeneity with respect to both surfaces. The scenarios do not permit early stopping for efficacy (i.e., $\delta_k = 0$ in the stopping rule), which is investigated in the next section. Early stopping for toxicity is still permitted, however, where we stop for toxicity in stratum k if $P(g(\widetilde{\mathbf{d}}) \leq g_k^\dagger \mid \mathcal{D}) < 0.9 \ \forall \ \widetilde{\mathbf{d}}$ is satisfied $J + 1 = 3$ consecutive times. In this event, the personalized algorithm permits dose-finding to continue in the other stratum unless it has also stopped for toxicity. We use values of Z_1 to index the true optimal dose combinations, $\mathbf{d}_{opt,k}$, and the values of the objective and toxicity functions at the $\mathbf{d}_{opt,k}$,

denoted by $f_{opt,k}$ and $g_{opt,k}$, respectively. To make the simulations more computationally feasible, we modify the algorithms to return a point estimate $\widehat{\mathbf{d}}_{opt,k} = \text{argmin}_{\widetilde{\mathbf{d}} \in \mathbb{A}_k} \mu(\widetilde{\mathbf{d}}, \widetilde{\mathbf{z}}_k)$ rather than the entire posterior distribution of $\mathbf{d}_{opt,k}$. Performance is compared for two values of the toxicity threshold $g_k^\dagger \in \{0.2, 0.5\}$ to assess how the distance of the $\mathbf{d}_{opt,k}$ from the $g_k^\dagger$ contour impacts the optimization. We do not consider a scenario where all dose combinations are toxic, since the dose ranges investigated in combination therapies are expected to be better targeted due to the previous early phase studies of each individual agent (Sweeting and Mander, 2012).

We utilize standardized dose combinations $\mathbf{d} = (d_1, d_2) \in [0, 1]^2$ which are separated in each coordinate direction by 0.25 units, creating a grid of 25 potential dose combinations in total. While the proposed method can optimize over the continuous dose combination space, engineering constraints may restrict the precision to which an agent can be manufactured. Thus, any constraints should be incorporated into the optimization, as we have done here. We set the point estimate of the recommended dose, $\widehat{\mathbf{d}}_{opt,k}$, as $\text{argmin}_{\widetilde{\mathbf{d}} \in \mathbb{A}_k} \mu(\widetilde{\mathbf{d}}, \widetilde{\mathbf{z}}_k)$ and the next dose combination to be evaluated, $\mathbf{d}^{(c_k+1)}$, as the maximizer of $\alpha_{cEI}(\widetilde{\mathbf{d}} \mid \mathcal{D}, \widetilde{\mathbf{z}}_k)$. For the initial data collection, we set $\rho_k = 0.25$ and utilize the modified $\mathbb{D}_k^{(q)*}$ to promote exploration at the earlier iterations of the algorithm. We define $\mathbb{D}_k^{(q)*}$ by setting $\mathbb{C}_k^{(q)} = \{\mathbf{d}_k^{(i)}\}_{i=0}^{q-1}$. Doing so prevents any previously evaluated dose combination from being selected as the next dose until after iteration $q = 8$, which is the iteration at which $\mathbb{D}_k$ becomes the unit hypercube. At each evaluated dose combination, the personalized algorithm uses $r_k = 2$ participants and the standard algorithm uses $r_k = 4$ participants. This ensures that both algorithms use the same number of total participants at each iteration, where the maximum sample size is set to be 80 participants.

We define $f(\mathbf{d}, z_1)$ and $g(\mathbf{d}, z_1)$ to be the continuous objective and toxicity surfaces which are defined using the multivariate normal density $h(\boldsymbol{\mu}) = N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with the same covariance matrix $\boldsymbol{\Sigma} = \left[\begin{smallmatrix} 0.1 & 0 \\ 0 & 0.1 \end{smallmatrix}\right]$ but with different mean vectors: $\boldsymbol{\mu}_1 = (0.5, 0.5)^T$, $\boldsymbol{\mu}_2 = (1, 1)^T$,

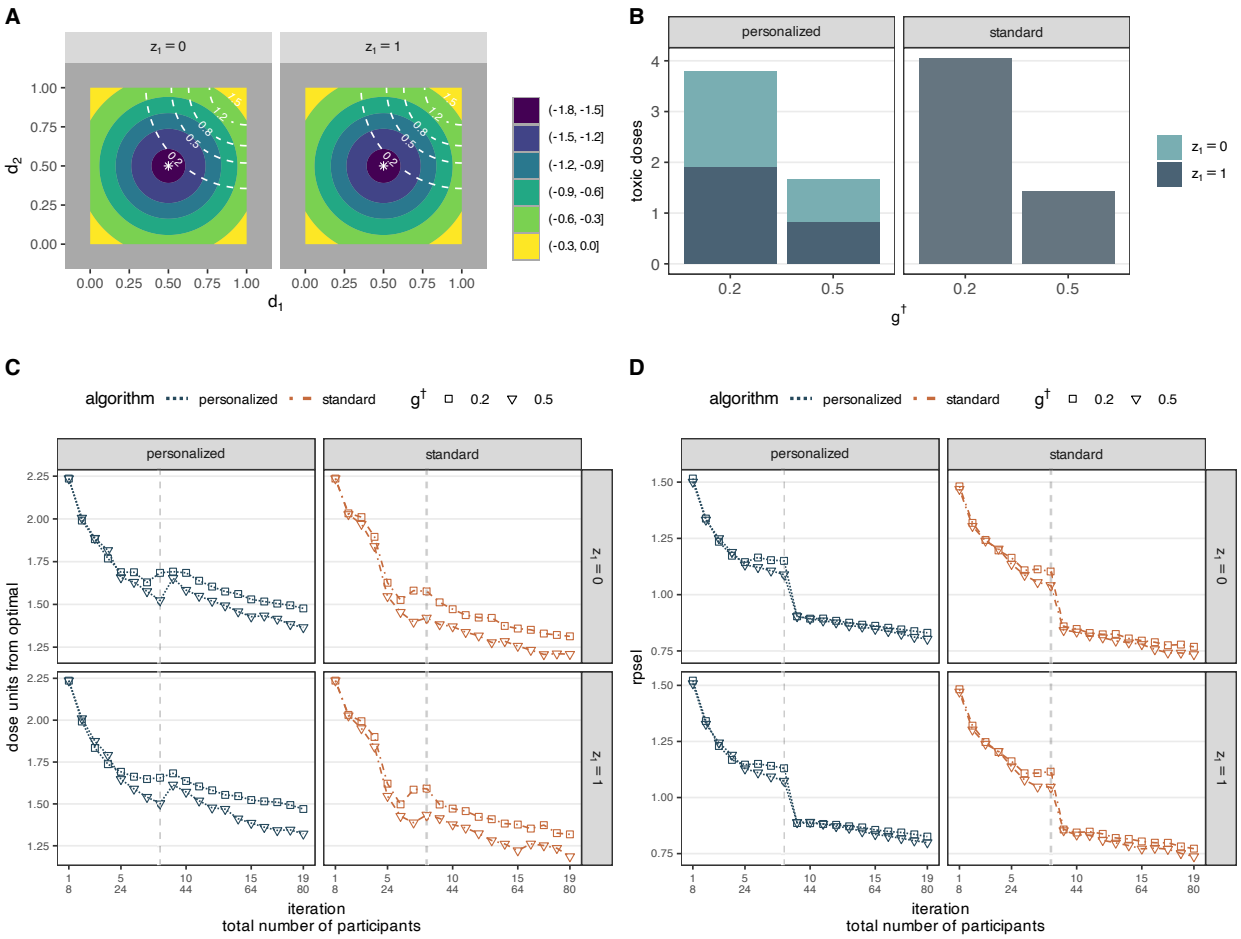

Figure 5.1: Scenario 1. A) Objective function, white stars denote $\mathbf{d}_{opt,k}$ and white dashed lines denote contours of toxicity function, B) expected number of truly toxic doses administered to participants as defined in (5.3), C) expected dose units from the optimal dose combination as defined in (5.4) by iteration, D) average RPSEL as defined in (5.5) by iteration. Vertical dashed lines in C and D represent iteration 8, after which the algorithm can search the entire dose combination space.

$\boldsymbol{\mu}_3 = (0.25, 0.75)^T$, $\boldsymbol{\mu}_4 = (0.75, 0.25)^T$, $\boldsymbol{\mu}_5 = (0.75, 1.25)^T$, and $\boldsymbol{\mu}_6 = (1.25, 0.75)^T$. For each scenario, the data generating mechanism is $y_f = f(\mathbf{d}, z_1) + \epsilon_f$ and $y_g = g(\mathbf{d}, z_1) + \epsilon_g$ where $\epsilon_f \sim N(0, \sigma^2_{y_f})$ and $\epsilon_g \sim N(0, \sigma^2_{y_g})$. The specifications of $f(\mathbf{d}, z_1)$ and $g(\mathbf{d}, z_1)$ are included in Table 5.1 (rows labeled "Simulation Study") and are plotted in panel A of Figures 5.1 and 5.2. The values of σ_{y_f} and σ_{y_g} are chosen to ensure specific standardized effect sizes (ses) at the optimal dose combination, defined as $ses_f = |f_{opt}|/\sigma_{y_f}$ and $ses_g = |g_{opt}|/\sigma_{y_g}$. Willard et al. (2023) considered several values for ses_f, which were 0.79/1/3.77 and represented the $25^{th}/50^{th}/75^{th}$ percentiles of ses from a meta-analysis of dose-responses for a large drug development portfolio at a pharmaceutical company (Thomas et al., 2014), and showed that a larger ses led to better performance of the algorithms. As we expect these results to generalize to the current setting, we focus specifically on $ses_f = ses_g = 1$, which represents the median ses we might expect for a new combination therapy. All computing is performed in R (R Core Team, 2022). The efficacy and toxicity functions are modeled using constant mean GP models with the anisotropic squared exponential kernel previously described. An empirical Bayes approach is adopted toward the GP hyperparameters (Binois and Gramacy, 2021), where initial values are utilized until convergence of the maximum likelihood estimation, which typically occurs within the first few iterations. For the lengthscale parameters, these values are $\sqrt{J+P}/2$, which is the midpoint of the maximal distance between two points on the unit cube defined by standardizing J dosing agents and P covariates. For the noise parameters, these values are the observed sample variances $\hat{\sigma}^2_f$ and $\hat{\sigma}^2_g$ for the objective and toxicity function models, respectively.

Three criteria are used to quantify algorithm performance and are estimated using $m = 1, ..., 1000$ Monte Carlo simulations. The expected number of toxic dose combinations administered to n_k participants is defined as follows:

$$E_{y|z_k}[\text{toxic doses}] \approx \frac{1}{1000} \sum_{m=1}^{1000} \left[\sum_{i=1}^{n_k} \mathbb{1}\left\{ g(\mathbf{d}_i^{(m)}, z_k) > g_k^\dagger \right\} \right]. \tag{5.3}$$

To assess how close $\widehat{\mathbf{d}}_k$ is to $\mathbf{d}_{opt,k}$ at each iteration, the expected number of dosing units is

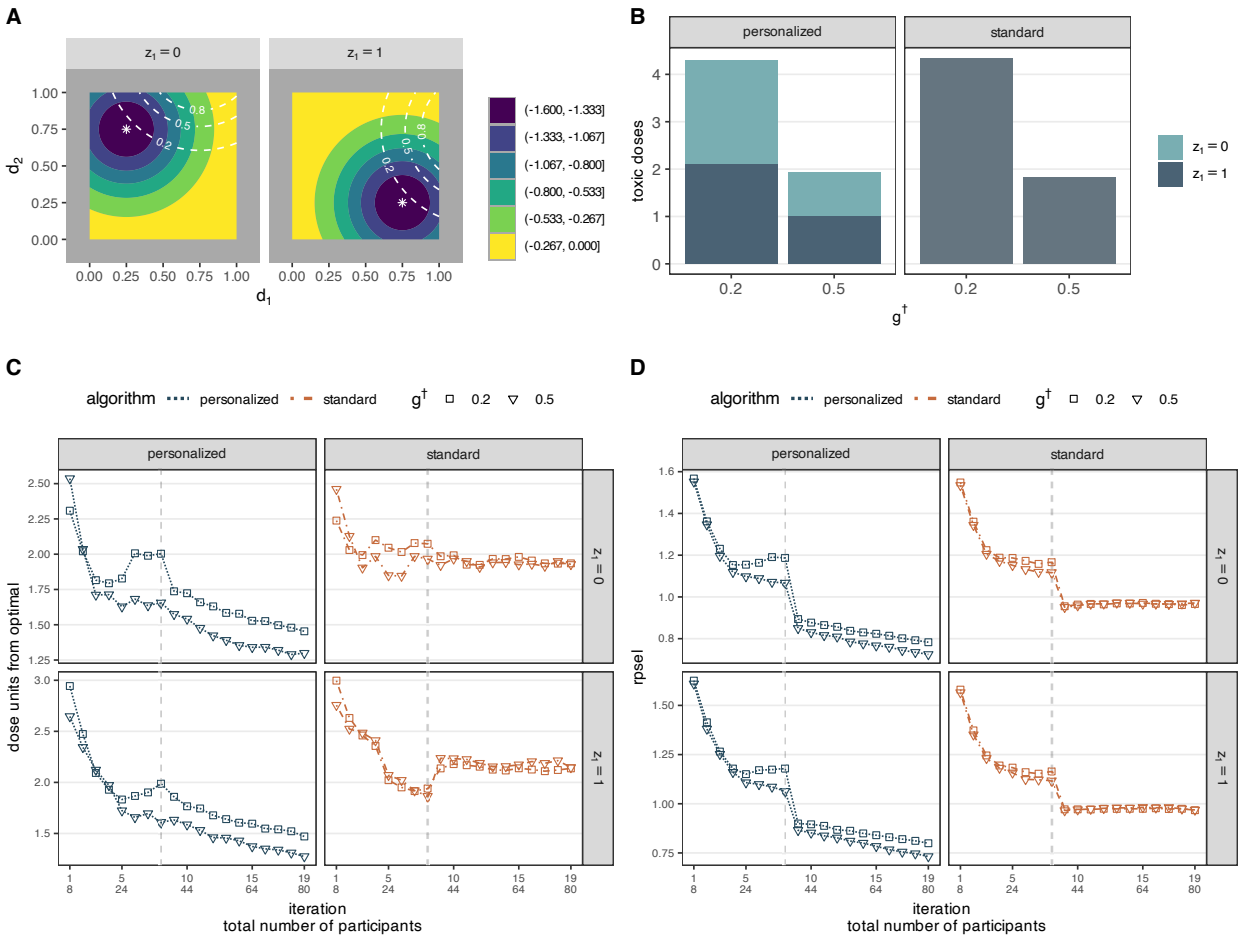

Figure 5.2: Scenario 2. A) Objective function, white stars denote $\mathbf{d}_{opt,k}$ and white dashed lines denote contours of toxicity function, B) expected number of truly toxic doses administered to participants as defined in (5.3), C) expected dose units from the optimal dose combination as defined in (5.4) by iteration, D) average RPSEL as defined in (5.5) by iteration. Vertical dashed lines in C and D represent iteration 8, after which the algorithm can search the entire dose combination space.

defined as their Euclidean distance divided by the distance between each dose combination, which is 0.25 in the simulations:

$$E_{y|z_k}[\text{dose units}] \approx \frac{1}{1000}\sum_{m=1}^{1000}\frac{\sqrt{(\hat{d}_1^{(m)} - d_{1,opt})_k^2 + (\hat{d}_2^{(m)} - d_{2,opt})_k^2}}{0.25}. \tag{5.4}$$

To assess how well the pointwise posterior distribution of the objective function at $\widehat{\mathbf{d}}_k$ estimates the true value $f_{opt,k}$, we use the expected root posterior squared error loss (RPSEL) which is estimated using $s = 1, ..., 10000$ posterior samples $f^{(s)}(\widehat{\mathbf{d}}_k, z_k)$:

$$E_{y|z_k}[RPSEL] \approx \frac{1}{1000}\sum_{m=1}^{1000}\left[\frac{1}{10000}\sum_{s=1}^{10000}(f^{(s)}(\widehat{\mathbf{d}}_k^{(m)}, z_k) - f_{opt,k})^2\right]^{\frac{1}{2}}. \tag{5.5}$$

Panels B-D of Figures 5.1 and 5.2 plot these criteria by iteration, showcasing the performance of the algorithms throughout the optimization process.

Scenario 1 considers the case of no response heterogeneity. Under this scenario, the personalized and standard algorithms administer a comparable number of toxic doses, where the number is higher when $\mathbf{d}_{opt,k}$ is closer to the tolerable toxicity contour $g_k^\dagger$ (Panel B of Figure 5.1). Under the more stringent toxicity threshold $g_k^\dagger = 0.2$, the personalized algorithm incorrectly stopped for toxicity in both strata 0.1% of the time, whereas the standard algorithm incorrectly stopped 0.3% of the time (results not shown). Neither algorithm incorrectly stopped under the more permissible toxicity threshold. A separate series of simulations was performed which did not perform dose escalation for the initial doses, but rather randomly selected the same number of initial doses as in the simulations described above (results not shown). Under this setting, the expected number of toxic doses administered was 11.9/11.9 for $g_k^\dagger = 0.2$ under the personalized/standard algorithm and 5.3/5.24 for $g_k^\dagger = 0.5$ under the personalized/standard algorithm. Thus, utilizing a dose escalation scheme results in a decrease in the number of toxic doses administered of roughly 66% and 74% for $g_k^\dagger = 0.2$ and $g_k^\dagger = 0.5$, respectively, underscoring the importance of dose escalation under higher-grade toxicity settings. Both the personalized and standard algorithms converge to the $\mathbf{d}_{opt,k}$,

with the optimization being more efficient when $\mathbf{d}_{opt,k}$ is further from the permissible toxicity contour (i.e., $g_k^\dagger = 0.5$ vs $g_k^\dagger = 0.2$ in Panel C of Figure 5.1). Additionally, they each yield a posterior distribution which estimates the $f_{opt,k}$ comparably (Panel D of Figure 5.1). The sudden drop in RPSEL after iteration 8 results from the dose escalation scheme ending. This permits the algorithms to explore any dose in the dose combination region which leads to improved optimization. Overall, both algorithms yield acceptable performance under no response heterogeneity, though, as expected, the standard algorithm is slightly more efficient.

Scenario 2 considers the case of response heterogeneity with respect to both the efficacy and toxicity surfaces. As in the previous scenario, the personalized and standard algorithms administer a comparable number of toxic doses, where the number is higher when $\mathbf{d}_{opt,k}$ is closer to the tolerable toxicity contour $g_k^\dagger$ (Panel B of Figure 5.2). Under the more stringent toxicity threshold $g_k^\dagger = 0.2$, the personalized algorithm incorrectly stopped for toxicity in both strata 0.1% of the time, whereas the standard algorithm incorrectly stopped more frequently at 1% of the time (results not shown). This difference results from the personalized algorithm permitting stratum-specific early stopping for toxicity, where the second stratum continues to be explored in the event of the first being stopped. Neither algorithm incorrectly stopped under the more permissible toxicity threshold. The personalized algorithm converges to the $\mathbf{d}_{opt,k}$, with the optimization being more efficient when $\mathbf{d}_{opt,k}$ is further from the permissible toxicity level (Panel C of Figure 5.1). However, the standard algorithm does not converge to $\mathbf{d}_{opt,k}$. As discussed in Willard et al. (2023), the marginal response surfaces are bimodal mixture distributions whose components consist of the equally weighted strata. Even though its use of the cEI may promote exploration of both modes during optimization, the standard algorithm still cannot determine how to optimally assign doses with respect to Z_1, a major disadvantage as compared to the personalized approach. Finally, the personalized algorithm yields posterior distributions which estimate the $f_{opt,k}$ well (Panel D of Figure 5.2), whereas the standard algorithm does not. Overall, the personalized algorithm is capable of

handling response heterogeneity with respect to the efficacy and toxicity surfaces, while the standard algorithm cannot.

5.5 Dose-finding design for obstructive sleep apnea therapy

Obstructive sleep apnea (OSA) is a common sleep disorder which affects over 930 million adults worldwide (Benjafield et al., 2019) and for which no pharmacotherapy has been approved. One continuous measure used to quantify OSA severity is the apnea-hypopnea index with 4% oxygen desaturation (AHI_4), which is measured in the number of apnea/hypopnea events per hour. Mild to moderate OSA is defined as 10-30 events per hour, and severe OSA is defined as 30 or more events per hour (Schweitzer et al., 2023). Based on a recently improved understanding of OSA's pathophysiology, several studies have proposed a combination of one of two antimuscarinic agents (oxybutynin and aroxybutynin) and a single norepinephrine reuptake inhibitor (atomoxetine) to serve as a potential pharmacotherapy (Schweitzer et al., 2023; Aishah et al., 2023; Rosenberg et al., 2022). These studies assessed several combinations of the dosing agents on the reduction in AHI_4 from baseline. The evaluation times ranged from a single night up to four weeks on treatment, and concluded that the proposed drug combinations display potential as a therapy for OSA. Furthermore, it was noted that a targeted therapeutic approach might prove beneficial in future studies to handle possible response heterogeneity (Aishah et al., 2023). The dose-finding design proposed below is inspired by findings from Schweitzer et al. (2023), which investigated the combination of aroxybutynin (0/2.5/5 mg) and atomoxetine (75 mg). While results were not analyzed separately by OSA severity subtype, Figure 2 of Schweitzer et al. (2023) suggests there is potential response heterogeneity with respect to these subtypes. For the mild to moderate subtype, the combinations (aroxybutynin/atomoxetine) 2.5mg/75mg and 5mg/75mg seem to yield a similar reduction in AHI_4, whereas for the severe subtype, the reduction under 5mg/75mg may be greater. However, it was observed that a larger number of AEs resulted

from the 5mg/75mg combination, though none were serious. To achieve maximal reduction in AHI_4 in the severe subtype, we may be willing to tolerate a larger number of AEs than for the mild to moderate subtype. In this case, the OBD combination would be 2.5mg/75mg for the mild to moderate subtype, but would be the 5mg/75mg combination for the severe subtype. Considering the above, we propose a personalized dose-finding design which tailors the OBD combination to OSA subtype and which permits different toxicity thresholds $g_k^\dagger$ for the severe versus mild to moderate subtypes.

To quantify the overall impact of AEs at each dose combination, we follow the approach in Le-Rademacher et al. (2020), who define a non-negative AE burden score that considers both the frequency and severity of the adverse events. Schweitzer et al. (2023) report the frequency of the most commonly occurring AEs; unfortunately, the grade (severity) information was not provided. However, since no severe adverse events were reported, we define the score by restricting our focus to grades 1-3 only, which we assume to occur 10%/45%/45% of the time independent of the type of AE, reflecting the belief that patients are more likely to remember and report higher graded events. Let Y_{icg} be the indicator that the i^{th} patient experiences an AE of type c which is of grade g, where $g = 1, 2, 3$ and where $c = 1, ..., 8$ corresponding to one of the following types: 1) dry mouth, 2) insomnia, 3) urinary hesitation/flow decrease, 4) constipation, 5) nausea, 6) decreased appetite, 7) feeling jittery, and 8) somnolence. Then patient i's AE burden is $B_i = \sum_c \sum_g w_{cg} Y_{icg}$. Specification of the weights w_{cg} is subjective and should be elicited from subject matter experts. We define $w_{cg} = w_c w_g$ and set grade weights $w_g = g$. Higher type weights w_c may be given to more burdensome types of AEs if desired. As was noted in (Schweitzer et al., 2023), the AEs which led to study discontinuation included insomnia, nausea, and dry mouth. We assign these AEs weights which are 5 times higher than the others and thus $\mathbf{w}_c = (5, 5, 1, 1, 5, 1, 1, 1)$. Interestingly, this study suggested that the 2.5mg/75mg combination decreased the frequency of AEs overall, as well as the frequency of more burdensome AEs as compared to the 0mg/75mg combination. This non-monotonitiicy in the dose-toxicity relationship is captured in the simulations that follow.

We define Z_1 as OSA severity, where $Z_1 = 0$ corresponds to the mild/moderate subtype and $Z_1 = 1$ corresponds to the severe subtype. We define $f(\mathbf{d}, z_1)$ as the reduction in AHI_4 from baseline and $g(\mathbf{d}, z_1) = \log(B + 0.5)$ as the log-transformed AE burden score. Using efficacy and adverse event results from (Schweitzer et al., 2023), we fit flexible linear models on standardized doses to determine the parameters used in the data generating mechanisms for $f(\mathbf{d}, z_1)$ and $g(\mathbf{d}, z_1)$, whose functional forms are included in Table 5.1 (rows labeled "OSA") and which are plotted in Panel A of Figure 5.3 on the original dosage scale. The $f(\mathbf{d}, z_1)$ differ across the strata, where the parameter vector for $Z_1 = 0$ is $\boldsymbol{\beta}_{Z_1=0} = (-1.38, -4.08, -0.48, -4.23, 2.45, -7.51, -1.56)$ and the parameter vector for $Z_1 = 1$ is $\boldsymbol{\beta}_{Z_1=1} = (1.05, -11.28, -8.32, -17.02, 8.17, 2.34, 4.61)$. The $g(\mathbf{d}, z_1)$ is the same across strata and has parameter vector $\boldsymbol{\theta} = (-0.59, 1.83, 2.26, -4.05, 1.79, 0.47, 2.91)$. As mentioned previously, under the personalized algorithm, we permit the tolerable toxicity thresholds to vary by strata. For the severe subtype, we set $g_1^\dagger = 2$, which equates to accepting an average AE burden score of about 7. This may include, for example, a single low-grade but high burden AE, or a couple moderate-grade but low burden AEs on average. For the mild/moderate subtype, we set $g_0^\dagger = 1.5$, which equates to accepting an average AE burden score of about 4. This may include, for example, a couple low-to-moderate-grade AEs of low burden, but would exclude higher burden AEs on average. The standard algorithm uses only a single value of the tolerable toxicity threshold, which we set equal to that of the mild/moderate subgroup, $g^\dagger = 1.5$.

Below we compare 12 designs using a maximum sample size of 100. The 12 designs are comprised of 6 personalized and 6 standard designs, where each uses a different combination of algorithm setting and stopping rule. The two settings differ in the number of patients assigned to each dose combination, where a larger number results in potentially fewer algorithm iterations, and a smaller number results in potentially more algorithm iterations. For the personalized algorithm, this is $r_k = 1$ and $r_k = 2$, and for the standard algorithm this is $r_k = 2$ and $r_k = 4$, respectively. We refer to these designs by their algorithm type and

the degree of replication: $P1/P2$ for the personalized designs, and $S2/S4$ for the standard designs. The three stopping rules investigated permit early stopping at roughly 40, 60, or 80 patients, and are denoted by $n_{STOP} = \{40, 60, 80\}$. Recall that we permit early stopping in stratum k when $\max_{\widetilde{\mathbf{d}} \in \mathbb{D}} \alpha_{cEI}(\widetilde{\mathbf{d}}_k \mid \mathcal{D}) < \delta_k$ at least $(J+1)$ times, which, for the dual-agent combination therapy being investigated, equates to 3 times. Within each design, we use the same value of δ_k for each stratum, and so $\delta_k = \delta$. The values of δ are calibrated such that each design achieves an expected sample size which is approximately the corresponding value of n_{STOP}. For $P1/P2$, $\delta = (0.1, 0.08, 0.065)/(0.137, 0.087, 0.07)$, and for $S2/S4$, $\delta = (0.077, 0.06, 0.05)/(0.1, 0.073, 0.055)$. The personalized algorithm permits stratum-specific early stopping, where the remaining budget is allocated to the other stratum where dose-finding may continue in the event of one of the strata stopping early. Finally, as there are no serious adverse events expected, we permit the modified dose combination region $\mathbb{D}_k^{(q)*}$ to expand at a quicker rate using $\rho_k = 0.5$. All other modeling and inferential details follow those previously described in the simulation study section. The performance of the designs is assessed and compared via the previously defined criteria which are estimated using 1,000 Monte Carlo replicates.

Under this scenario, we find that all designs yield expected sample sizes corresponding to their respective values of n_{STOP} (Panel B of Figure 5.3). The personalized algorithms converge to the $\mathbf{d}_{opt,k}$, with the optimization being more efficient in stratum $Z_1 = 0$ which has the larger ses_f (Panel D of Figure 5.3). The standard algorithms are more efficient in converging to $\mathbf{d}_{opt,0}$, but fail to converge to $\mathbf{d}_{opt,1}$. Similar results hold for estimation of $f_{opt,k}$, where the personalized algorithms can handle the heterogeneity across strata whereas the standard algorithms cannot. We note that the personalized designs generally lead to a slightly larger number of participants receiving toxic doses (Panel C of Figure 5.3), but this results from these designs exploring more doses on average. For example, designs $P1/P2$ under $n_{STOP} = 40$ are expected to evaluate 11/9 unique doses whereas designs $S2/S4$ under $n_{STOP} = 40$ are expected to evaluate 7/6 unique doses (results not shown). We find that a lower degree of

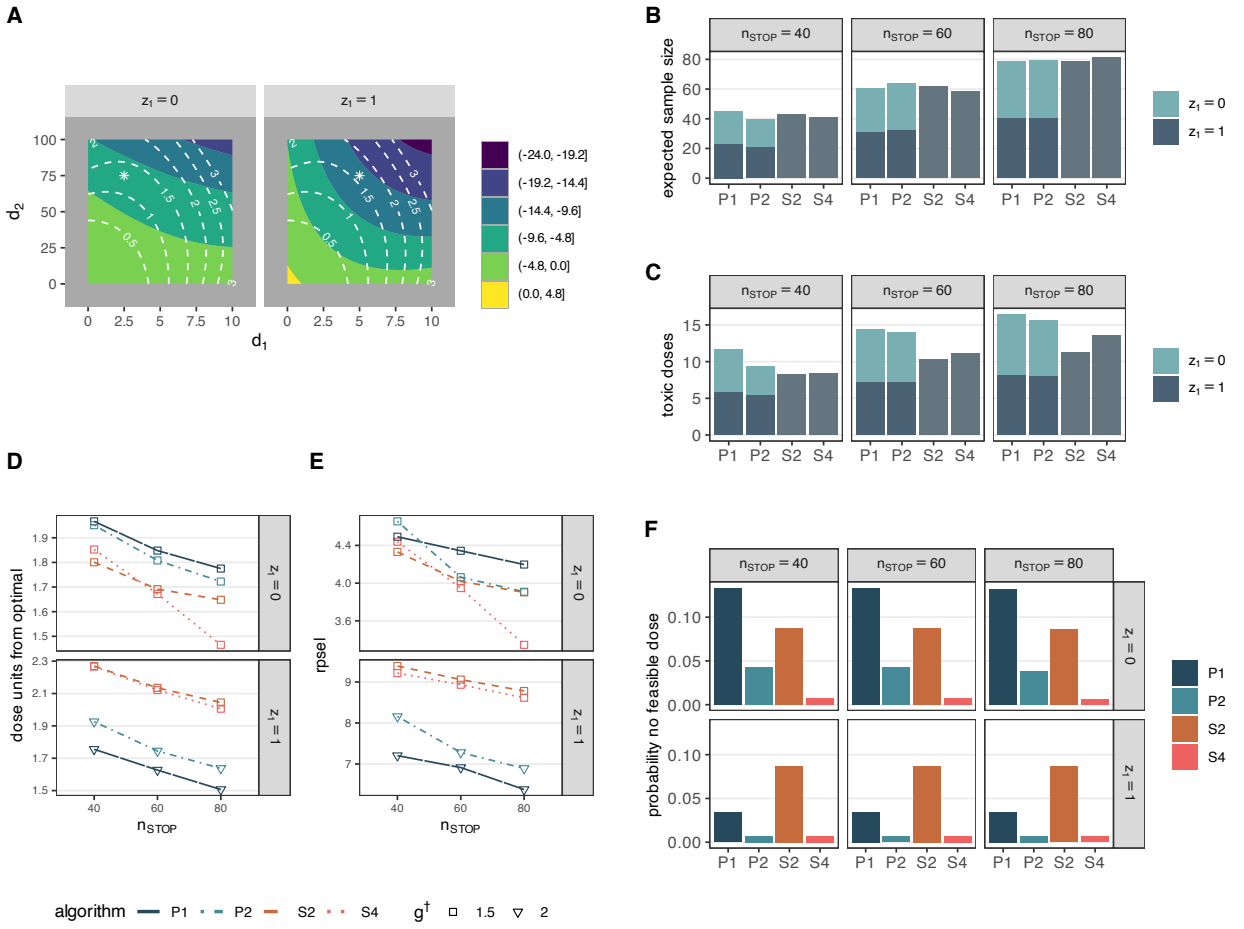

Figure 5.3: A) Objective function, white stars denote $\mathbf{d}_{opt,k}$, white dashed lines denote contours of toxicity function, B) expected sample size, C) expected number of truly toxic doses administered to participants as defined in (5.3), D) expected dose units from the optimal dose combination as defined in (5.4) by iteration, E) average RPSEL as defined in (5.5) by iteration, and F) probability of incorrectly stopping for toxicity and declaring no feasible doses. For OSA subtype, $Z_1 = 0$ denotes mild/moderate and $Z_1 = 1$ denotes severe.

replication may lead to an increased probability of incorrectly determining that no feasible doses exist (i.e., $P1/S2$ vs $P2/S4$ in Panel F of Figure 5.3), and that for the personalized algorithm, this is more extreme in stratum $Z_1 = 0$ which has the more stringent toxicity threshold $g_0^\dagger = 1.5$. Note that the standard algorithm yields a single decision only, so this probability is fixed across strata. Given the goal of the trial is to identify subtype-tailored dose combinations, the standard designs would be unusable in practice. Furthermore, we exclude design $P1$ from consideration as it yields a substantial increase in the probability of declaring no feasible doses as compared to $P2$. Based upon the simulated scenario, we would recommend $P2$ with either $n_{STOP} = 60$ or $n_{STOP} = 80$ given their improvements over $P2$ with $n_{STOP} = 40$. Design $P2$ with $n_{STOP} = 60$ is expected to evaluate 10 unique doses in 64 patients with 14 receiving toxic doses on average, whereas design $P2$ with $n_{STOP} = 80$ is expected to evaluate 11 unique doses in 80 patients with 16 receiving toxic doses on average. The study investigator could make a final decision by performing a cost-benefit analysis using these estimates.

5.6 Discussion

In this work, we generalized the methods proposed in Willard et al. (2023) to the setting of higher-grade continuous toxicities under no monotonicity assumptions. We showed how these toxicities could be incorporated into the search strategy employed by the Bayesian optimization methods, as well as proposed a dose-escalation scheme to collect the initial data. A personalized approach is shown to be beneficial when response heterogeneity exists in at least one of the efficacy and toxicity surfaces. Additionally, the personalized approach permits tolerable levels of toxicity to be defined with respect to individual strata, providing an additional level of tailoring which is not possible when utilizing the standard approach.

The proposed work is not without limitations. Firstly, all simulations were performed using a single binary covariate under the setting of dual-agent dose combinations only. Under the personalized approach, the use of additional categorical variables or dosing agents is

straightforward, though would require larger sample sizes. Extension to continuous covariates remains as future work. While conceptually straightforward under the proposed approach, the dangers and consequences of inadvertently extrapolating into regions of covariate combinations never before seen in the trial would need to be considered and assessed. Secondly, future research should investigate the impact of a wider range of dose escalation schemes on overall algorithm performance than those that are featured here. As this may be applicable to a variety of constrained Bayesian optimization problems outside of dose-finding, this is a particularly interesting research direction. Finally, we assumed conditional independence between the efficacy and toxicity surfaces. In trials with binary efficacy and toxicity responses, Cunanan and Koopmeiners (2014) showed that estimating correlations can be challenging given the small sample sizes of early phase designs and that designs which assumed independence between efficacy and toxicity can still perform well. Further investigations for continuous responses are warranted, however. To relax the assumption of conditional independence, the efficacy and toxicity responses could be jointly modeled using a multivariate GP. The cEI acquisition function would need to be updated to incorporate the correlation between the two responses. We leave this as a direction for future work.

Changes to the GP modeling and acquisition function may help improve performance of the algorithms. Stationary, anisotropic squared exponential kernels were used but could be replaced by *non-stationary* kernels which include dose-covariate interaction terms. These would permit the correlation between two dose combinations to depend on the actual dosage levels and strata in which they lie, rather than solely on their distance from one another in the input space. Doing so, however, may greatly increase the dimensionality of the hyperparameter space, potentially leading to the same estimation challenges the proposed methods sought to avoid. Additionally, an empirical Bayes approach was adopted toward the GP hyperparameters, where initial values of the hyperparameters were used until the maximum likelihood estimation routines converged. While convenient and practical, proceeding in a fully Bayesian manner by placing priors on all hyperparameters may improve estimation at

the earlier stages of optimization. However, this comes at the increased computational cost of posterior sampling, which could be prohibitive depending on the number of scenarios of interest and the expected number of algorithm iterations, each of which requires refitting the GP models. Finally, the cEI acquisition function utilizes a point estimate of the current best observation, $f_k^* = \min_{\widetilde{\mathbf{d}} \in \mathbb{A}_k} \boldsymbol{\mu}_f(\widetilde{\mathbf{d}}_k)$, which ignores the uncertainty around this quantity. Proceeding in a fully Bayesian manner by integrating over its posterior distribution may yield efficiency gains. Two different acquisition functions adopt this approach (Gramacy et al., 2016; Letham et al., 2019), but are not available analytically and require Monte Carlo for evaluation. While utilization of these acquisition functions would greatly increase the computational burden of the proposed dose-finding algorithm, they represent an exciting direction for future work.

Data availability

No new data were created or analyzed in this study. The R scripts used for the simulations and graphics can be found on a public GitHub repository at https://github.com/jjwillard/bayesopt_obdc.

Acknowledgments

The authors would like to thank Dr. John Kimoff for his valued input regarding the OSA application component of this manuscript. JW acknowledges the support of a doctoral training scholarship from the Fonds de recherche du Québec - Nature et technologies (FRQNT). SG acknowledges the support by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada (NSERC), a Chercheurs-boursiers (Junior 1) award from the Fonds de recherche du Québec - Santé (FRQS), and support from the Canadian Statistical Sciences Institute. EEMM acknowledges support from a Discovery Grant from NSERC. EEMM is a Canada Research Chair (Tier 1) in Statistical Methods for Precision Medicine and acknowledges the support of a chercheur de mérite career award from the FRQS. This

research was enabled in part by support provided by Calcul Québec and the Digital Research Alliance of Canada.

Chapter 6

Conclusion

6.1 Summary

In this thesis, I demonstrated how utilizing covariate information benefits both early phase and late phase Bayesian adaptive clinical trials. In the first manuscript introduced in Chapter 3, I described how covariate adjustment has been shown to increase the statistical power of fixed size RCTs. I investigated the impact of combining covariate adjustment with prior information on the covariate effects in Bayesian adaptive designs with early stopping criteria. I showed how covariate adjustment leads to trials which have greater power, and which stop earlier and more often. These findings hold for trials with a variety of endpoints and sample sizes, as well as for both collapsible and non-collapsible estimands. I found there to be no added benefit to using informative priors over weakly informative ones when the effective sample sizes were moderate or larger, though I noted that informative priors may be beneficial in situations where this is not the case (e.g., trials investigating rare diseases). Furthermore, I found that inclusion of non-prognostic variables may have deleterious effects in trials that use binary endpoints and have small maximum sample sizes, and so caution must be exercised under this scenario. The work presented in Chapter 3 provides further

theoretical support for the FDA's recently published draft guidance on covariate adjustment in RCTs (United States Food and Drug Administration, 2021). Since it is already common to collect a substantial amount of covariate information within the scope of a trial (Friedman et al., 2015), covariate adjustment may provide great benefit for little additional cost.

Through the second and third manuscripts presented in Chapters 4 and 5, respectively, I described the goal of personalized dose-finding trials for combination therapies as identifying patient-specific optimal dose combinations. I emphasized the challenge of extending many standard parametric methods to the personalized setting since they require potentially many dose-covariate interaction terms which may be poorly estimated given the small sample sizes in these trials. These estimation challenges are exacerbated in the combination therapy setting. I showed how using Bayesian optimization makes personalized dose-finding for combination therapies feasible. Furthermore, under response heterogeneity, I found the personalized approach to be essential, since the standard "one-size-fits-all" approach recommends only a single dose which may fail to be optimal for some, or even all, patients in the population.

In Chapter 4, focus was placed on a minimal toxicity setting, where toxicity information was not considered during the dose-finding. Estimating the efficacy function using a non-parametric GP model achieved the flexibility required to model non-monotonicities and avoided the estimation challenges that would potentially arise in similarly flexible, but highly parameterized, models. By pairing this with an acquisition function which utilized stratum-specific information, patient-specific optimal doses could be obtained. This approach was generalized to the setting of higher-grade toxicities in Chapter 5, where toxicity information was incorporated into the dose-optimization to identify patient-specific optimal biological dose combinations. The sequential search strategy was defined using an acquisition function which balanced the patient-specific exploration-exploitation trade-off for efficacy, while giving higher weights to safe dose combinations. A dose escalation scheme was proposed to collect

the initial responses while at the same time minimize the number of patients given toxic dose combinations. This was shown to provide a large reduction in the number of toxic dose combinations which were administered as compared to a standard space filling design which employed random sampling. From a trial ethics perspective, this is an especially important contribution. Finally, the personalized approach allowed for different levels of permissible toxicity to be defined with respect to disease severity, providing an even greater degree of tailoring over standard dose-finding approaches. Overall, the proposed methods provide potential avenues toward achieving improved dose-optimization, a primary goal of the FDA's newly created Project Optimus (United States Food and Drug Administration, 2023).

6.2 Limitations and future work

The proposed methods are not without limitations. Below I discuss areas for improvement and future research, focusing first on the work proposed for late phase Bayesian adaptive RCTs and then proceeding to the methods proposed for early phase Bayesian adaptive personalized dose-finding trials.

6.2.1 Late phase Bayesian adaptive RCTs

The covariate adjustment performed in Chapter 3 assumed joint independence between the adjustment variables. This represents an ideal scenario which may not apply in all situations. If the variables were instead correlated, they would carry similar information and we would expect that adjusting for them would result in smaller increases in power and the probability of stopping early than what is observed under independence. Thus, fewer independent covariates would be needed to obtain the same benefit as a set of correlated covariates. This helps motivate the question of the appropriate number of covariates to include for adjustment. A related issue is the required strength of association between a covariate and outcome to consider it for adjustment. A limitation of the work proposed in Chapter 3 is that only a single set of covariate strengths was assessed for each endpoint. However, other simulations were performed (but not reported) which showed that covariates

which were only weakly associated with the outcome yielded minimal benefit. This finding agrees with observations from Cox and McCullagh (1982) where, for a normal endpoint and single continuous adjustment covariate, they note that a correlation of less than 0.3 yields minimal increase in treatment effect precision whereas a correlation greater than 0.9 yields a substantial increase in precision. Furthermore, they add that when the number of degrees of freedom in the trial is small, only one or two adjustment variables should be used.

The question of the appropriate number of covariates to adjust for is an important one and should be addressed by future work. These investigations could be motivated by the variable selection and shrinkage literature, or the model selection and averaging literature. For example, non-parametric approaches, such as GP models, could be used for covariate adjustment. While these may be more robust than parametric approaches in settings where the functional relationship between the outcome and the adjustment covariates is difficult to correctly specify, they are computationally intensive and so were not considered in this work. Furthermore, if multiple adjustment models were under consideration, Bayesian model averaging could be employed. Under this approach, parameter estimates are obtained by taking a weighted average of model-specific parameter estimates, with the weights corresponding to the posterior probability of each model being true given the data. This might guard against model misspecification, but comes at an increased computational cost and so was not investigated in the current work. While it is possible that employing more flexible or data-driven approaches for covariate adjustment provides benefit, these approaches are often deleterious if not specified in an objective manner and so were not investigated in this thesis. For example, if interim model selection procedures favor the alternative hypothesis, an inflation of Type I error rate is expected (Pocock et al., 2002). The controlling and estimation of operating characteristics should therefore account for model selection as well as for efficacy decisions. Future work could investigate the impact of data-driven approaches toward covariate adjustment.

The covariate adjustment performed in Chapter 3 was in the context of an adaptive *stopping* rule, where simple randomization was performed. Additional benefit might be observed in the same trial by also utilizing adjustment in an adaptive *allocation* rule. Recall that adaptive allocation rules allow for changes to the randomization probabilities throughout the course of the trial. Covariate-adjusted response adaptive randomization updates the allocation probability to a particular treatment arm based on previously observed responses for a particular covariate pattern (e.g., Rosenberger and Sverdlov (2008)). As this type of adaptive randomization may lead to covariate imbalance across treatment arms, additionally adjusting for the covariates in the adaptive stopping rule should be beneficial. In designs which use response adaptive randomization, the allocation rules generally dictate randomizing a patient to a particular treatment arm in proportion to its posterior probability of being the superior treatment arm, and often include additional tuning parameters (e.g., Thall and Wathen (2007)). When covariates are included, the allocation rule has a greater chance of randomizing patients to better performing treatment arms with respect to their covariate pattern, which implicitly defines a goal of assigning patients to their patient-specific optimal treatment. This goal is the same as that of the personalized dose-finding methods proposed in Chapters 4 and 5, and so we might be interested in adapting these methods to define an allocation rule. For example, optimizing an acquisition function defines a deterministic allocation rule, where patients are assigned to the treatment arm which is the maximizer of the patient-specific acquisition function. A randomized allocation scheme could be performed by defining patient-specific treatment arm randomization probabilities as the normalized values of a patient-specific acquisition function. As many acquisition functions naturally accommodate the trade-off between exploration and exploitation of the objective function and are grounded in statistical decision theory, their use within randomized allocation rules could be highly beneficial. This is a very interesting direction for future work and combines elements from the methods proposed in Chapters 3-5.

Another limitation of the methods in Chapter 3 is that covariate adjustment was performed

using outcome regression models only. Inverse probability of treatment weighting is another method which is used to obtain estimates of marginal treatment effects and which has been discussed within the frequentist clinical trial literature (see Morris et al. (2022)). Inverse probability of treatment weighting uses estimates of patient-specific propensity scores to re-weight the trial's sample such that the distribution of covariates is balanced across treatment arms. When the marginal treatment effect contrast is formed using the re-weighted samples from each treatment arm, an unbiased estimate of the marginal treatment effect is obtained. Saarela et al. (2015b) show how this weighting may be performed through the Bayesian paradigm, which provides one potential avenue toward investigating its use in Bayesian adaptive clinical trials. However, weighting can be less efficient, especially for smaller sample sizes, so this approach was not investigated in this thesis though remains an exciting direction for future work.

Finally, the methods utilized throughout Chapter 3 assumed the magnitudes of the covariate effects were known. In practice, these are likely to be unknown and so basing power and sample size calculations on them may be undesirable. Benkeser et al. (2021) suggest using an unadjusted analysis to conservatively estimate power and required sample sizes, with any estimates provided by a covariate adjusted analysis being reported as potential efficiency gains under a best-case scenario only. This thesis sought only to show the potential benefit of covariate adjustment and so did not investigate this practical limitation. Future work should consider covariate adjustment approaches which do not rely upon assumed covariate effect sizes.

6.2.2 Early phase Bayesian adaptive dose-finding trials

In this subsection, I discuss limitations of the personalized dose-finding approach proposed in Chapters 4 and 5, noting that many of the same limitations apply under both the minimal and higher-grade toxicity settings and so they will be discussed together. To begin, the personalized dose-finding approach was illustrated using discrete covariates only. Extension to

continuous covariates would be highly valuable, since covariates which are thought to have non-linear relationships with an outcome (e.g., age) should not be discretized. While incorporating continuous covariates into the proposed method is conceptually straightforward, extrapolating into unobserved regions of the covariate space is an important concern. Recall that the stratum-specific acquisition function conditions on a particular covariate pattern and assigns the next dose combination as its maximizer. When using continuous covariates, a dose assignment is likely to be made for a patient whose covariate pattern has never been observed, even near the end of the trial. Under this scenario, the acquisition function is still a well-defined functional of the posterior distribution at a specific input (i.e., dose combination and covariate pattern), but as a result of potentially sparse data in that region, the posterior may be either very diffuse or simply maintain the features of the prior distribution. Future work should investigate this and propose appropriate guidelines for recommending OBD combinations in this context.

The proposed personalized dose-finding approach focuses on tailoring OBD combinations to patient-specific covariate profiles. There may arise situations where interest lies in tailoring with respect to a specific set of covariates while also adjusting for a different set of prognostic covariates. We saw in Chapter 3 how adjusting for prognostic covariates can increase the precision of the estimated treatment effect and the power of randomized controlled trials. Following a similar line of reasoning for personalized dose-finding, we might hope to adjust for variables known to be associated with the outcome but which are not of interest for tailoring. This may increase the precision of the estimated dose response and improve dose-finding overall. To accomplish this, a GP model using the full set of covariates could be fitted after which standardization could be performed with respect to the adjustment variables only. This would serve to marginalize over the adjustment variables, resulting in estimated patient-specific responses (i.e., defined by the tailoring covariates) which may have greater precision than their unadjusted counterparts. While the limited sample sizes in early phase dose-finding trials might make it challenging to accommodate both tailoring and adjustment

variables, this remains an interesting direction for future work and combines themes from Chapters 3-5.

Another limitation of the proposed method is the assumption of conditional independence between the efficacy and toxicity functions given dose. When this assumption does not hold, the correlation between the responses must be modeled. This can be done via a multivariate GP model. Under this scenario, a fully Bayesian scheme could be advantageous as the small sample sizes might make the correlation parameter challenging to estimate, and so using a more informative prior distribution for the correlation parameter could help the model to converge. As previously mentioned, the work proposed in this thesis did not perform personalized dose-finding in a fully Bayesian manner since many common posterior approximation methods are too computationally demanding in this setting. For example, Markov chain Monte Carlo is one such approximation method which requires expensive posterior sampling to be performed at each iteration of the dose-finding. Furthermore, it is inefficient in a sequential setting since it completely discards samples from previous iterations. A sequential Monte Carlo approach may prove to be a better alternative here, as it is able to update, rather than discard, information about the posterior distribution at each iteration of dose-finding (e.g., Gramacy and Polson (2011); Golchi and Thorlund (2020)). To facilitate fully Bayesian inference in the proposed dose-finding algorithms, future work should investigate more efficient sequential computational schemes such as this.

The personalized dose-finding approaches proposed under the minimal and higher-grade toxicity settings each used a variant of the Expected Improvement acquisition function. Under the minimal toxicity setting, the Augmented Expected Improvement (Huang et al., 2006) was used, and under the higher-grade toxicity setting, the constrained Expected Improvement (Gardner et al., 2014) was used. One limitation of having used these variants of Expected Improvement is that the uncertainty around the best observed efficacy function value f_k^* was ignored. That is, a plug-in estimate for f_k^* was used rather than considering its posterior

distribution. To fully account for the uncertainty in f_k^* in a constrained optimization context, Letham et al. (2019) proposed the Noisy Expected Improvement. This acquisition function handles uncertainty in a fully Bayesian way by integrating over the posterior distributions of any parameters not of direct interest. Furthermore, it permits batch optimization, where more than a single dose combination may be explored at each iteration. This may be useful in contexts which have long delays between administration of the dose and observation of the response. While incorporating the Noisy Expected Improvement into the dose-finding procedure is an exciting direction for future work, the acquisition function must be evaluated through Monte Carlo simulation. This will tremendously increase the computational requirements of the dose-finding algorithm and requires development of efficient computational tools.

Finally, the personalized dose-finding methods were assessed under simulation scenarios determined by the authors of the manuscripts included in Chapters 4 and 5. While all efforts were made to include a broad and realistic set of scenarios under which the proposed methods were assessed, there may exist some scenarios which were not considered. Additionally, the simulations assumed that the covariate patterns represented by each strata were equally prevalent and had the same enrollment rates throughout the trials, and so an equal number of participants per strata were enrolled at each iteration of the personalized dose-finding algorithm. This represents a best-case scenario in terms of power and the ability to enroll, but in practice it is improbable that all subgroups are equally represented in the target population. In the extreme scenario where one strata has zero patients, personalized dose-finding for this stratum could not proceed, as any recommendation made by the algorithm would simply reflect the prior information used at the start of the trial. Furthermore, due to information borrowing across strata, the estimation for sparse strata may be dominated by strata which have many patients. For this reason, Zhang et al. (2024) suggest subgroup-specific dose-optimization should only be performed when pre-defined subgroups enroll enough patients in the trial, a recommendation which supports the simulation scenarios evaluated in

the present work.

6.3 Concluding remarks

In this thesis, I demonstrated that covariate adjustment improves the operating characteristics of late-phase Bayesian adaptive randomized controlled trials. Additionally, I developed methods for personalized adaptive dose-finding trials for combination therapies under minimal and higher-grade toxicity settings. Through three manuscripts, I showed that incorporating additional covariate information into the decision making process used throughout Bayesian adaptive clinical trials offers great benefit and does so for potentially little cost. Furthermore, a tutorial on posterior marginalization and algorithms for dose-finding were provided in an effort to make the methods easier to implement in practice. Therefore, this thesis promotes adoption of advanced methodology in clinical trials to improve their efficiency and probability of success.

Appendices

APPENDIX A

Appendix to Manuscript 1

A.1 Marginalization procedures

A.1.1 Continuous outcome: difference in means

Let Y be a Normally distributed outcome with the target of inference being the marginal difference in means:

$$\gamma(\boldsymbol{\theta}) := \mu(\boldsymbol{\theta}; A = 1) - \mu(\boldsymbol{\theta}; A = 0)$$

where $\mu(\boldsymbol{\theta}; A = a) = E[Y \mid A = a; \boldsymbol{\theta}]$. Under the assumption of at least one treatment-covariate interaction (i.e., $\mathbf{Z} \neq \emptyset$; treatment effect heterogeneity), the difference in means is non-collapsible. Estimation proceeds assuming independent outcomes and the following model:

$$p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta}) = \text{Normal}(\mu(\boldsymbol{\theta}; A_i, \mathbf{X}_i), \sigma^2)$$

$$\mu(\boldsymbol{\theta}; A_i, \mathbf{X}_i) = \beta_0 + \phi A_i + \mathbf{X_i}\boldsymbol{\beta} + (A_i \cdot \mathbf{Z_i})\boldsymbol{\omega}$$

$$\boldsymbol{\theta} = \{\beta_0, \phi, \boldsymbol{\beta}, \boldsymbol{\omega}, \sigma^2\}$$

$$\mathcal{L} = \prod_{i=1}^{n_t} p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta})$$

$$\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t}) \propto \prod_{i=1}^{n_t} p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta}) p(\boldsymbol{\theta}).$$

Samples from the posterior distribution of $\gamma(\boldsymbol{\theta})$ are obtained from adjusted analyses by marginalizing $s = 1, ..., S$ samples of the conditional $\mu(\boldsymbol{\theta}; A, \mathbf{X})$ before forming the contrast:

$$\begin{aligned}\gamma(\boldsymbol{\theta}_s) &= \mu(\boldsymbol{\theta}_s; A = 1) - \mu(\boldsymbol{\theta}_s; A = 0) \\ &= \int_{\mathbf{X}} \mu(\boldsymbol{\theta}_s; A = 1, \mathbf{X}) p(\mathbf{X}) d\mathbf{X} - \int_{\mathbf{X}} \mu(\boldsymbol{\theta}_s; A = 0, \mathbf{X}) p(\mathbf{X}) d\mathbf{X} \\ &= \int_{\mathbf{X}} (\beta_{0,s} + \phi_s + \mathbf{X}\boldsymbol{\beta}_s + \mathbf{Z}\boldsymbol{\omega}_s) p(\mathbf{X}) d\mathbf{X} - \int_{\mathbf{X}} (\beta_{0,s} + \mathbf{X}\boldsymbol{\beta}_s) p(\mathbf{X}) d\mathbf{X}.\end{aligned}$$

The integrals are approximated using the Bayesian bootstrap procedure described in Section 3.3 of the manuscript. After fitting the linear model, $s = 1, ..., S$ samples are obtained from the joint posterior distribution of the model parameters, $\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t})$. Let $\boldsymbol{\theta}_s$ represent the s^{th} draw from this joint posterior distribution. For every row $i = 1, ..., n_t$ in the sample data, a value of $A_i = 1$ is assigned. Then for each $\boldsymbol{\theta}_s$ the following procedure is performed. The n_t values of the linear predictor $\mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{x}_i)$ are calculated. A vector $\mathbf{w}_s = (w_{1,s}, ..., w_{n_t,s})$ is drawn from a Dirichlet($\mathbf{1}_{n_t}$) distribution. Using $\mathbf{w}_s$, the n_t values are then averaged, $\sum_{i=1}^{n_t} w_{i,s} \mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{x}_i)$, which marginalizes them with respect to the observed $\mathbf{X} = \mathbf{x}$, yielding a single sample $\mu(\boldsymbol{\theta}_s; A = 1)$. This occurs for all $\boldsymbol{\theta}_s$ to yield S samples from the posterior distribution of $\mu(\boldsymbol{\theta}; A = 1)$. This entire process is then repeated for $A_i = 0$, to yield S samples from the posterior distribution of $\mu(\boldsymbol{\theta}; A = 0)$. These posterior samples are then subtracted to yield samples from the posterior distribution of $\gamma(\boldsymbol{\theta})$. A brief summary outline is included below.

1. Fit the linear regression model with identity link.
2. Obtain $s = 1, ..., S$ samples from the joint posterior distribution of the model parameters, $\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t})$.
3. Create one copy of the sample data where $A_i = 1$ for all $i = 1, ..., n_t$.
4. For each $\boldsymbol{\theta}_s$, perform the following:
   (a) For each $i = 1, ..., n_t$, calculate $\mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{x}_i)$.

(b) Sample $\mathbf{w}_s = (w_{1,s}, ..., w_{n_t,s})$ from a Dirichlet($\mathbf{1}_{n_t}$) distribution.

(c) Average these n_t values to marginalize with respect to the observed $\mathbf{X} = \mathbf{x}$:

$\mu(\boldsymbol{\theta}_s; A_i = 1) = \sum_{i=1}^{n_t} w_{i,s}\mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{X}_i)$.

5. The $s = 1, ..., S$ values of $\mu(\boldsymbol{\theta}_s; A = 1)$ are samples from the posterior of $\mu(\boldsymbol{\theta}; A = 1)$.

6. Repeat steps 3-4 for $A_i = 0$ to yield S samples from the posterior of $\mu(\boldsymbol{\theta}; A = 0)$.

7. Subtract to obtain S samples from the posterior of $\gamma(\boldsymbol{\theta})$.

A.1.2 Binary outcome: relative risk

Letting Y be distributed as a Bernoulli random variable, where $Y = 1$ indicates an event occurs and $Y = 0$ indicates no event occurs, a marginal estimand of interest is the relative risk:

$$\gamma(\boldsymbol{\theta}) := \mu(\boldsymbol{\theta}; A = 1)/\mu(\boldsymbol{\theta}; A = 0)$$

where $\mu(\boldsymbol{\theta}; A = a) = E[Y \mid A = a; \boldsymbol{\theta}]$. Estimation proceeds assuming independent outcomes and the following model:

$$p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta}) = \text{Bernoulli}(\mu(\boldsymbol{\theta}; A_i, \mathbf{X}_i))$$
$$\mu(\boldsymbol{\theta}; A_i, \mathbf{X}_i) = \text{logit}^{-1}(\beta_0 + \phi A_i + \mathbf{X}_i\boldsymbol{\beta} + (A_i \cdot \mathbf{Z}_i)\boldsymbol{\omega})$$
$$\boldsymbol{\theta} = \{\beta_0, \phi, \boldsymbol{\beta}, \boldsymbol{\omega}\}$$

$$\mathcal{L} = \prod_{i=1}^{n_t} p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta})$$
$$\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t}) \propto \prod_{i=1}^{n_t} p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta})p(\boldsymbol{\theta}).$$

To obtain posterior samples of the marginal estimand from adjusted analyses, $s = 1, ..., S$ posterior samples from the inverse logit link function applied to the linear predictors under treatment and no treatment are marginalized and then divided:

$$\begin{aligned}
\gamma(\boldsymbol{\theta}_s) &= \frac{\mu(\boldsymbol{\theta}_s; A = 1)}{\mu(\boldsymbol{\theta}_s; A = 0)} \\
&= \frac{\int_{\mathbf{X}} \mu(\boldsymbol{\theta}_s; A = 1, \mathbf{X})p(\mathbf{X})d(\mathbf{X})}{\int_{\mathbf{X}} \mu(\boldsymbol{\theta}_s; A = 0, \mathbf{X})p(\mathbf{X})d(\mathbf{X})} \\
&= \frac{\int_{\mathbf{X}} \text{logit}^{-1}\{\beta_{0,s} + \phi_s + \mathbf{X}\boldsymbol{\beta}_s + \mathbf{Z}\boldsymbol{\omega}_s\}p(\mathbf{X})d(\mathbf{X})}{\int_{\mathbf{X}} \text{logit}^{-1}\{\beta_{0,s} + \mathbf{X}\boldsymbol{\beta}_s\}p(\mathbf{X})d(\mathbf{X})}
\end{aligned}$$

The integrals are approximated using the Bayesian bootstrap procedure described in Section 3.3 of the manuscript. After fitting the generalized linear model with logit link function corresponding to η, $s = 1, ..., S$ samples are obtained from the joint posterior distribution of the model parameters, $\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t})$. Let $\boldsymbol{\theta}_s$ represent the s^{th} draw from this joint posterior distribution. For every row $i = 1, ..., n_t$ in the sample data, a value of $A_i = 1$ is assigned. Then for each $\boldsymbol{\theta}_s$ the following procedure is performed. The n_t values of the indexed linear predictors are calculated and transformed by the inverse logit to yield samples from $\mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{x}_i) = \text{logit}^{-1}(\beta_{0,s} + \phi_s A_i + \mathbf{X}_i\boldsymbol{\beta}_s + (A_i \cdot \mathbf{Z}_i)\boldsymbol{\omega}_s)$. A vector $\mathbf{w}_s = (w_{1,s}, ..., w_{n_t,s})$ is drawn from a Dirichlet($\mathbf{1}_{n_t}$) distribution. Using $\mathbf{w}_s$, the n_t values are then averaged, $\sum_{i=1}^{n_t} w_{i,s}\mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{x}_i)$, which marginalizes them with respect to the observed $\mathbf{X} = \mathbf{x}$, yielding a single sample $\mu(\boldsymbol{\theta}_s; A = 1)$ from the posterior distribution of $\mu(\boldsymbol{\theta}; A = 1)$. This occurs for all $\boldsymbol{\theta}_s$ to yield S samples from the posterior distribution of $\mu(\boldsymbol{\theta}; A = 1)$. This entire process is then repeated for $A_i = 0$, to yield S samples from the posterior distribution of $\mu(\boldsymbol{\theta}; A = 0)$. These are then divided to yield S samples from the posterior distribution of $\gamma(\boldsymbol{\theta})$. A brief summary outline is included below.

1. Fit the logistic regression model.
2. Obtain $s = 1, ..., S$ samples from the joint posterior distribution of the model parameters, $\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t})$.
3. Create one copy of the sample data where $A_i = 1$ for all $i = 1, ..., n_t$.
4. For each $\boldsymbol{\theta}_s$, perform the following:
   (a) For each $i = 1, ..., n_t$, calculate $\mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{x}_i) = \text{logit}^{-1}(\beta_{0,s} + \phi_s A_i + \mathbf{X}_i\boldsymbol{\beta}_s + (A_i \cdot \mathbf{Z}_i)\boldsymbol{\omega}_s)$.
   (b) Sample $\mathbf{w}_s = (w_{1,s}, ..., w_{n_t,s})$ from a Dirichlet($\mathbf{1}_{n_t}$) distribution.
   (c) Average these n_t values to marginalize with respect to the observed $\mathbf{X} = \mathbf{x}$: $\mu(\boldsymbol{\theta}_s; A = 1) = \sum_{i=1}^{n_t} w_{i,s}\mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{x}_i)$.

5. The $s = 1, ..., S$ values of $\mu(\boldsymbol{\theta}_s; A = 1)$ are samples from the posterior distribution of $\mu(\boldsymbol{\theta}; A = 1)$.

6. Repeat steps 3-4 for $A_i = 0$ to yield S samples from the posterior distribution of $\mu(\boldsymbol{\theta}; A = 0)$.

7. Divide $\mu(\boldsymbol{\theta}_s; A = 1)/\mu(\boldsymbol{\theta}_s; A = 0)$ for each s to obtain S samples from the posterior distribution of $\gamma(\boldsymbol{\theta})$.

A.1.3 Time-to-event outcome: hazard ratio

Let $Y = \{T, \delta\}$ be defined as in the section for hazard ratios in the manuscript, where the target of inference is the marginal hazard ratio:

$$\begin{aligned}\gamma(\boldsymbol{\theta}) &= h(t \mid A = 1)/h(t \mid A = 0) \\ &= \log\{\mu(\boldsymbol{\theta}; A = 1)\}/\log\{\mu(\boldsymbol{\theta}; A = 0)\}\end{aligned}$$

where $\mu(\boldsymbol{\theta}; A = a) = S(t \mid A = a; \boldsymbol{\theta})$. Estimation proceeds assuming independent outcomes, no competing risks, and the following model:

$$h_i(t \mid A_i, \mathbf{X}_i) = h_0(t)\exp(\eta_i)$$

$$\eta_i = \phi A_i + \mathbf{X}_i\boldsymbol{\beta} + (A_i \cdot \mathbf{Z_i})\boldsymbol{\omega}$$

$$S_i(t \mid A_i, \mathbf{X}_i) = \exp\left(-I(t; \boldsymbol{\psi}, \boldsymbol{k}, \delta)\exp(\eta_i)\right)$$

$$\boldsymbol{\theta} = \{\boldsymbol{\psi}, \phi, \boldsymbol{\beta}, \boldsymbol{\omega}\}$$

$$p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta}) = S_i(T_i \mid A_i, \mathbf{X}_i)^{1-\delta_i} h_i(T_i \mid A_i, \mathbf{X}_i)^{\delta_i}$$

$$\mathcal{L} = \prod_{i=1}^{n_t} p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta})$$

$$\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t}) \propto \prod_{i=1}^{n_t} p(Y_i \mid A_i, \mathbf{X}_i, \boldsymbol{\theta}) p(\boldsymbol{\theta}).$$

As the hazard ratio is non-collapsible, $s = 1, ..., S$ marginal posterior samples from adjusted analyses are obtained through marginalization of the log transformed survival probabili-

ties:

$$\begin{aligned}\gamma(\boldsymbol{\theta}_s) &= \frac{h_s(t \mid A=1)}{h_s(t \mid A=0)} \\ &= \frac{\log\{\mu(\boldsymbol{\theta}_s; A=1)\}}{\log\{\mu(\boldsymbol{\theta}_s; A=0)\}} \\ &= \frac{\log\{\int_{\mathbf{X}} \mu(\boldsymbol{\theta}_s; A=1, \mathbf{X}) p(\mathbf{X}) d\mathbf{X}\}}{\log\{\int_{\mathbf{X}} \mu(\boldsymbol{\theta}_s; A=0, \mathbf{X}) p(\mathbf{X}) d\mathbf{X}\}} \\ &= \frac{\log\{\int_{\mathbf{X}} \exp\left[-I(t; \boldsymbol{\psi}_s, \boldsymbol{\kappa}, \delta) \exp(\phi_s + \mathbf{X}\boldsymbol{\beta}_s + \mathbf{Z}\boldsymbol{\omega}_s)\right] p(\mathbf{X}) d\mathbf{X}\}}{\log\{\int_{\mathbf{X}} \exp\left[-I(t; \boldsymbol{\psi}_s, \boldsymbol{\kappa}, \delta) \exp(\mathbf{X}\boldsymbol{\beta}_s)\right] p(\mathbf{X}) d\mathbf{X}\}}.\end{aligned}$$

Dividing log-transformed survival probabilities can be numerically unstable and is undefined for all t such that $S(t|A=a) \in \{0,1\}$. Thus we use a more numerically stable identity (Stitelman et al., 2011; Remiro-Azócar et al., 2022):

$$\begin{aligned}\log(\gamma(\boldsymbol{\theta}_s)) &= \log\left(\frac{h_s(t|A=1)}{h_s(t|A=0)}\right) = \log\left(\frac{-\log[\mu(\boldsymbol{\theta}_s; A=1)]}{-\log[\mu(\boldsymbol{\theta}_s; A=0)]}\right) \\ &= \log\{-\log[\mu(\boldsymbol{\theta}_s; A=1)]\} - \log\{-\log[\mu(\boldsymbol{\theta}_s; A=0)]\}.\end{aligned}$$

The integrals are approximated using the Bayesian bootstrap procedure described in Section 3.3 of the manuscript. After fitting the flexible, semi-parametric proportional hazards model, $s = 1, ..., S$ samples are obtained from the joint posterior distribution of the model parameters, $\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t})$. Let $\boldsymbol{\theta}_s$ represent the s^{th} draw from this joint posterior distribution. For every row $i = 1, ..., n_t$ in the sample data, a value of $A_i = 1$ is assigned. Then for each $\boldsymbol{\theta}_s$ the following procedure is performed. For the t corresponding to the time from the start of the trial to the current analysis, the n_t values of the indexed conditional survival probabilities, $\mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{x}_i)$ are calculated. A vector $\mathbf{w}_s = (w_{1,s}, ..., w_{n_t,s})$ is drawn from a Dirichlet($\mathbf{1}_{n_t}$) distribution. Using $\mathbf{w}_s$, the n_t values are then averaged, $\sum_{i=1}^{n_t} w_{i,s}\mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{x}_i)$, which marginalizes them with respect to the observed $\mathbf{X} = \mathbf{x}$, yielding a single sample $\mu(\boldsymbol{\theta}_s; A = 1)$ from the posterior distribution of $\mu(\boldsymbol{\theta}; A = 1)$. For numerical stability, a $\log\{-\log[\cdot]\}$ transformation is applied to yield a single sample from the posterior distribution of $\log\{h(t \mid A = 1)\}$. This occurs for all $\boldsymbol{\theta}_s$ to yield S draws from the posterior distribution of $\log\{h(t \mid A = 1)\}$. This entire process is then repeated

for $A_i = 0$, to yield S draws from the posterior distribution of $\log\{h(t \mid A = 0)\}$. These posterior draws are subtracted and then exponentiated to yield samples from the posterior distribution of the marginal hazard ratio $\gamma(\boldsymbol{\theta})$. A brief summary is below.

1. Fit a flexible semi-parametric proportional hazards model.
2. Obtain $s = 1, ..., S$ samples from the joint posterior distribution of the model parameters, $\pi(\boldsymbol{\theta} \mid \mathcal{D}_{n_t})$.
3. Create one copy of the sample data where $A_i = 1$ for all $i = 1, ..., n_t$.
4. For each $\boldsymbol{\theta}_s$, perform the following:
   (a) For each $i = 1, ..., n_t$, calculate the conditional survival probabilities at time t corresponding to the time from the start of the trial to the current analysis, $\mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{x}_i)$.
   (b) Sample $\mathbf{w}_s = (w_{1,s}, ..., w_{n_t,s})$ from a Dirichlet$(\mathbf{1}_{n_t})$ distribution.
   (c) Average these n_t values to marginalize with respect to the observed $\mathbf{X} = \mathbf{x}$: $\mu(\boldsymbol{\theta}_s; A = 1) = \sum_{i=1}^{n_t} w_i \mu(\boldsymbol{\theta}_s; A_i = 1, \mathbf{X}_i = \mathbf{x}_i)$.
   (d) Apply a $\log\{-\log[\cdot]\}$ transformation to yield a single sample from the posterior distribution of $\log\{h(t \mid A = 1)\}$.
5. This yields S samples from the posterior distribution of $\log\{h(t \mid A = 1)\}$.
6. Repeat steps 3-4 for $A_i = 0$ to yield S samples from the posterior distribution of $\log\{h(t \mid A = 0)\}$.
7. Subtract and then exponentiate to obtain S samples from the posterior distribution of the marginal hazard ratio, $\gamma(\boldsymbol{\theta})$.

A.2 Ascertainment of marginal estimand values

A.2.1 Ascertainment of marginal relative risk (binary outcome)

We first select a value of β_0 on the log-odds scale in the adjusted data generating models, such that the simulated datasets have a specific marginal control event risk (p_{ctr}). We then use β_0 and the obtained values of ϕ (i.e., those where the unadjusted model achieves 50% and 80% power, also on the log-odds scale) to select the reported value of the marginal relative risk.

To find β_0, let Y be a binary outcome. As a reminder, we define A as the treatment assignment indicator, where $A = 1$ means being assigned to the treatment group and $A = 0$ means being assigned to the control group. Let l be the number of participants assigned to control, k be the number of participants assigned to treatment, and $l + k = n$ be the total number of participants potentially enrolled in the trial. Let $\mathbf{X}_{n\times p}$ be the set of covariates used in the adjusted data generating model, and $\mathbf{X}_i$ be the row vector corresponding to the values of the covariates for the i^{th} participant. Recall the marginal control event risk, p_{ctr}, is the risk of having an event in those assigned to control. Then p_{ctr} can be defined with respect to an adjusted data generating model as follows:

$$
\begin{aligned}
p_{ctr} &= E[Y|A=0] \\
&\approx \frac{1}{l}\sum_{i=1}^{l} \hat{E}[Y_i|A_i=0,\mathbf{X}_i] \\
&= \frac{1}{l}\sum_{i=1}^{l} \text{logit}^{-1}\{\beta_0 + \phi(A_i=0) + \mathbf{X}_i\boldsymbol{\beta}\} \\
&= \frac{1}{l}\sum_{i=1}^{l} \text{logit}^{-1}\{\beta_0 + \mathbf{X}_i\boldsymbol{\beta}\} \\
0 &= \left[\sum_{i=1}^{l} \text{logit}^{-1}\{\beta_0 + \mathbf{X}_i\boldsymbol{\beta}\}\right] - l \times p_{ctr}
\end{aligned}
$$

Given a fixed value of p_{ctr}, conditional covariate effects $\boldsymbol{\beta}$ on the log-odds scale, and initial simulation of the treatment assignment and covariate distributions, $\{A, \mathbf{X}\}$, β_0 can be optimized using the last line above (i.e., using uniroot() in R).

In the simulations for each relative risk value within each maximum sample size, 5,000 datasets (each with 5,000 participants) were generated using $\{\boldsymbol{\beta}, A, \mathbf{X}\}$ as described under the binary outcome data generating mechanism. From these, 5,000 values for β_0 were found and the mean of this distribution was selected as the value of β_0. Using this and the value of ϕ, 5,000 values for the marginal relative risk were obtained by dividing the proportion of events in those assigned to treatment ($\hat{E}[Y|A=1]$) by the proportion of events in those assigned to control ($\hat{E}[Y|A=0]$). The mean of this distribution was then reported as the value of the marginal relative risk corresponding to β_0 and ϕ.

A.2.2 Ascertainment of marginal hazard ratio (time-to-event outcome)

Our goal is specify a value of the reported marginal hazard ratio which corresponds to the value of ϕ (on the log-hazard scale) used in the adjusted data generating models. Recall our assumption of proportional hazards, where the marginal hazard ratio is not time-dependent. Let $Y = \{T, \delta\}$ be defined as in the section for hazard ratios in the manuscript. Define A as the treatment assignment indicator, where $A = 1$ means being assigned to the treatment group and $A = 0$ means being assigned to the control group. Let t be the maximum duration of the trial and $P(T > t|A = 1) = S(t|A = 1)$ and $P(T > t|A = 0) = S(t|A = 0)$ be the survival probabilities at time t for those assigned to treatment and control, respectively. In the simulations for each hazard ratio value within each maximum sample size, 5,000 datasets (each with 5,000 participants) were generated using $\{\boldsymbol{\beta}, A, \mathbf{X}, t = 50\}$ as described under the time-to-event outcome data generating mechanism. For each dataset, the value of

the marginal hazard ratio was calculated as:

$$\gamma = \exp\{\log(-\log[\hat{P}(T > 50|A = 1)]) - \log(-\log[\hat{P}(T > 50|A = 0)])\}$$

The mean of this distribution was then reported as the value of the marginal hazard ratio corresponding to ϕ.

A.3 CCEDRN COVID-19 RCT truncated covariate distributions

The following algorithm yields datasets from truncated Normal distribution F with correct post-truncation minimum $\wedge$, maximum $\vee$, $q1$, and $q3$ values and approximately correct post-truncation μ and σ values.

1. Determine maximum sample size of trial, max_ss.
2. Obtain reported summary statistics $\{q1, q3, \mu, \sigma\}$ and range $[\wedge, \vee]$.
3. Set plausible value of median $q2$ if not reported.
4. Set $n = 250 \times max\_ss$.
5. Select starting values for $\{\xi, \tau^2\}$.
6. Until $\mu^* \approx \mu$ and $\sigma^* \approx \sigma$:
   (a) Generate $i = 1, ..., n$ values of X_i from $N(\xi, \tau^2)$.
   (b) Discard $X_i \notin [min, max]$.
   (c) Sample $\frac{n}{4}$ values of $X_i \in [\wedge, q1]$.
   (d) Sample $\frac{n}{4}$ values of $X_i \in [q1, q2]$.
   (e) Sample $\frac{n}{4}$ values of $X_i \in [q2, q3]$.
   (f) Sample $\frac{n}{4}$ values of $X_i \in [q3, \vee]$.
   (g) Collect all and set μ^* and σ^* as the mean and standard deviation of the sampled values.
   (h) Update $\{\xi, \tau^2\}$ or break if $\mu^* \approx \mu$ and $\sigma^* \approx \sigma$.
7. To generate one dataset from F, use final values of $\{\xi, \tau^2\}$ to repeat process above,

but sample max_ss values from the n collected values.

For age, the following summary statistics and simulation parameter values were used:

$$\{\wedge = 18, q1 = 39, q2 = 55, q3 = 70, \vee = 90, \mu = 54.7, \sigma = 19.8, \xi = 62, \tau = 40\}.$$

For respiratory rate, the following summary statistics and simulation parameter values were used:

$$\{\wedge = 12, q1 = 18, q2 = 20, q3 = 22, \vee = 40, \mu = 21.0, \sigma = 6.2, \xi = 30, \tau = 6\}.$$

A.4 Summary graphics for bias and RMSE

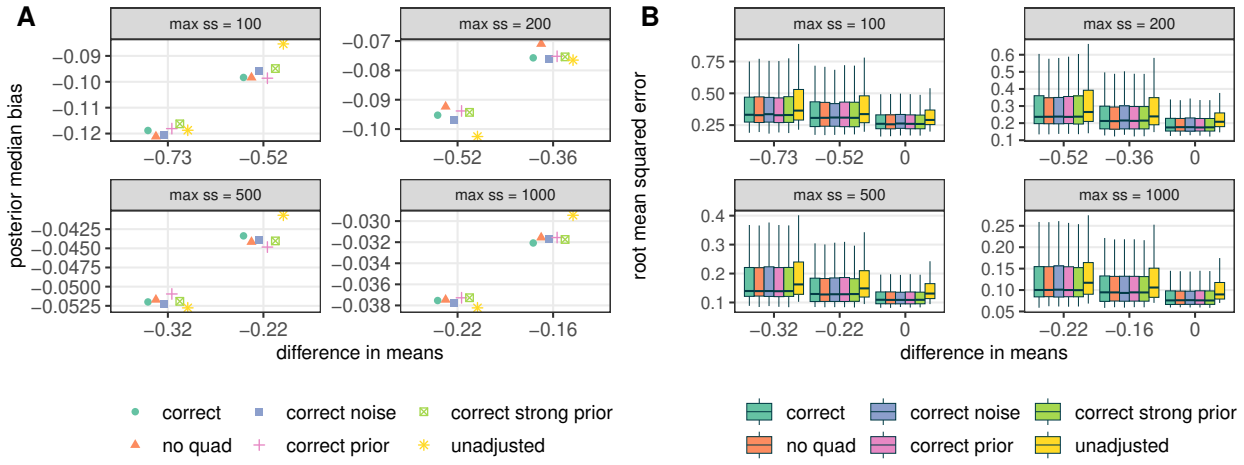

Figure A.4.1: Continuous outcome. A) Posterior median bias and B) root mean squared error. Panels correspond to various maximum sample sizes (max ss). Points are jittered horizontally.

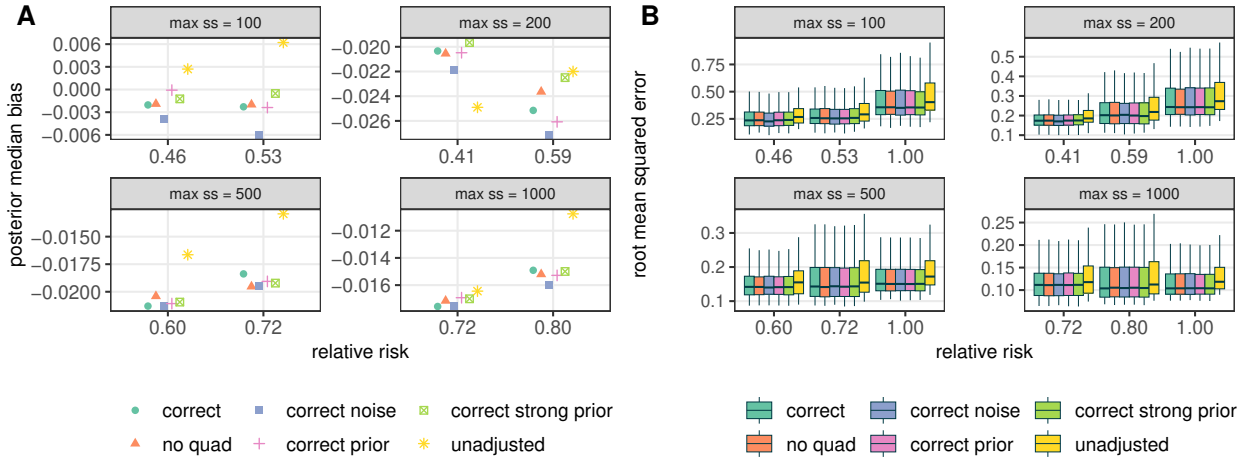

Figure A.4.2: Binary outcome. A) Posterior median bias and B) root mean squared error. Panels correspond to various maximum sample sizes (max ss). Points are jittered horizontally.

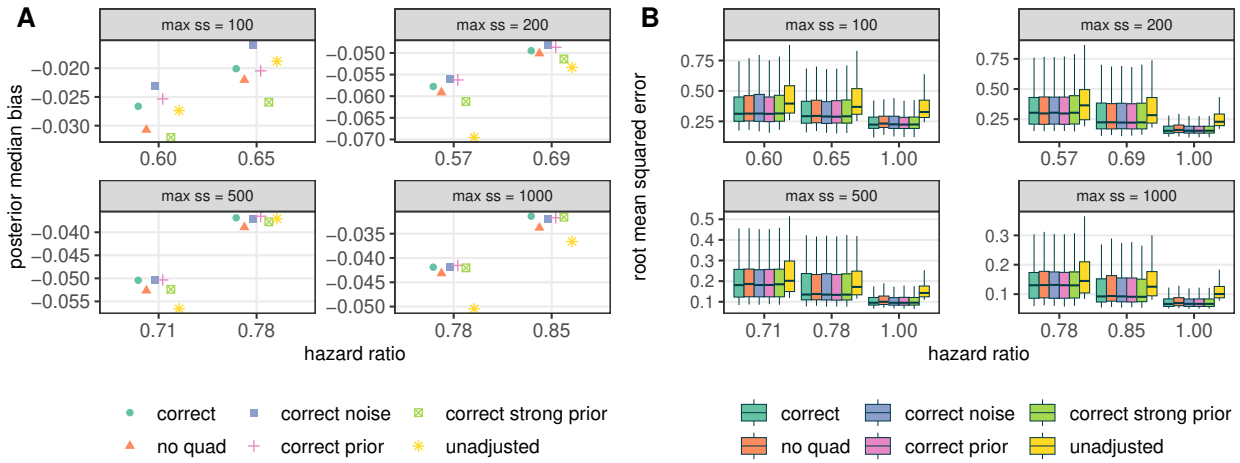

Figure A.4.3: Time-to-event outcome. A) Posterior median bias and B) root mean squared error. Panels correspond to various maximum sample sizes (max ss). Points are jittered horizontally.

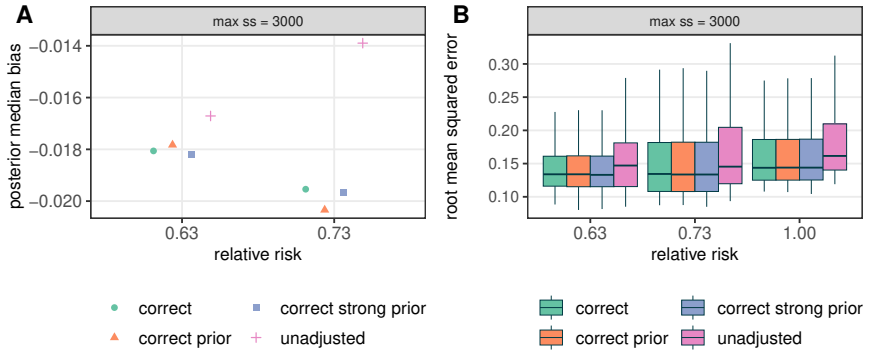

Figure A.4.4: CCEDRN-ADAPT COVID-19 trial with binary outcome. A) Posterior median bias and B) root mean squared error. Panels correspond to a maximum sample size (max ss) of 3,000. Points are jittered horizontally.

A.5 Simulation for informative prior on treatment effect

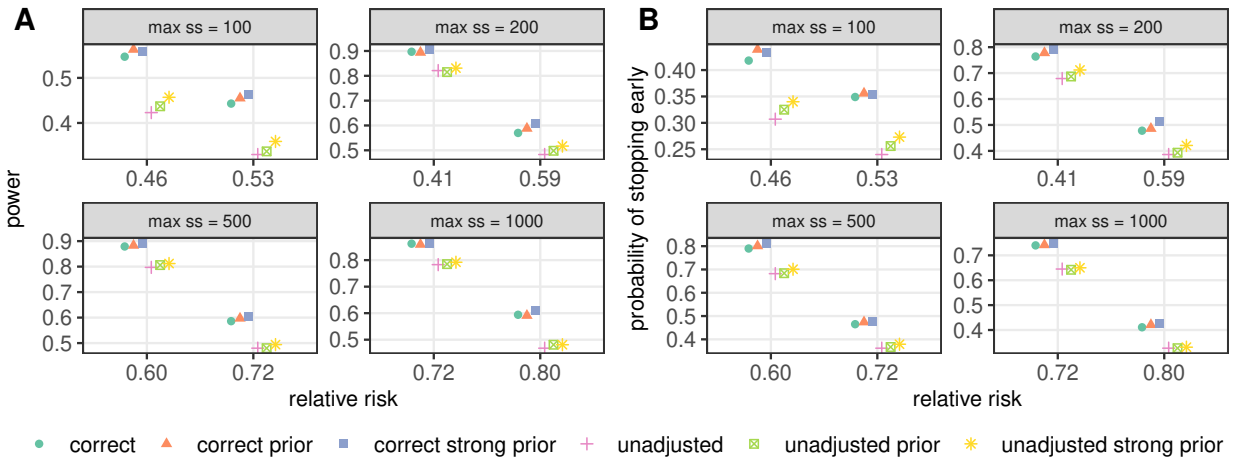

Figure A.5.1: Binary outcome. A) Power and B) probability of stopping early. Panels correspond to various maximum sample sizes (max ss). Points are jittered horizontally.

In this section, we consider the impact of incorporating informative prior information on the treatment effect for trials with binary endpoints. Six adjustment models are considered:

1. correct: $\beta_0 + \phi A + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_3^2 + \beta_5 X_5$
2. correct prior: $\beta_0 + \phi A + \beta_1^\dagger X_1 + \beta_2^\dagger X_2 + \beta_3^\dagger X_3 + \beta_4^\dagger X_3^2 + \beta_5^\dagger X_5$
3. correct strong prior: $\beta_0 + \phi A + \beta_1^{\dagger\dagger} X_1 + \beta_2^{\dagger\dagger} X_2 + \beta_3^{\dagger\dagger} X_3 + \beta_4^{\dagger\dagger} X_3^2 + \beta_5^{\dagger\dagger} X_5$
4. unadjusted: $\beta_0 + \phi A$
5. unadjusted prior: $\beta_0 + \phi^\dagger A$
6. unadjusted strong prior: $\beta_0 + \phi^{\dagger\dagger} A$

The regression coefficients $\{\phi, \boldsymbol{\beta}, \boldsymbol{\beta}^\dagger, \boldsymbol{\beta}^{\dagger\dagger}\}$ and the functional forms of the *correct*, *correct prior*, *correct strong prior*, and *unadjusted* models are defined as previously described for the binary trials in the simulation study section of the manuscript. The *unadjusted prior* model includes a prior on the treatment indicator coefficient centered at the value used in the data

generating mechanism where the *unadjusted* model achieves approximately 80% power (50% power for max ss = 100). The *unadjusted strong prior* model both centers and re-scales this prior to be more informative. The following priors are then used for the treatment indicator coefficients in the *unadjusted prior* and *unadjusted strong prior* models:

$$\phi^{\dagger} \sim \text{Normal}(c, 2.5/s_a)$$

$$\phi^{\dagger\dagger} \sim \text{Normal}(c, 1/s_a)$$

where $c = \{-1.21, -1.36, -0.82, -0.54\}$ for max ss $= \{100, 200, 500, 1000\}$. All other components for the binary trials remain as described in the simulation study section of the manuscript.

Results for power and the probability of stopping early are displayed in Figure A.5.1. Including stronger priors on the treatment effect may increase the power and probability of stopping early as compared to weakly informative priors. This hold for both the *correct* and *unadjusted* model variants and is most beneficial for smaller sample sizes. However, this comes at the cost of inflated type 1 error (T1E), with the greatest inflation occurring for the smaller maximum sample sizes (Table A.5.1). Both the type 1 error inflation and increase in power become less pronounced in the trials with larger maximum sample sizes where the priors are dominated by the data.

Table A.5.1: Binary outcome. Type 1 error rate (T1E), bias under the null (Bias*), and expected sample size at three different values of the marginal relative risk (γ).

	Maximum sample size = 100					Maximum sample size = 200				
			Expected sample size					Expected sample size		
Adjustment model	T1E	Bias*	$\gamma = 1$	$\gamma = 0.53$	$\gamma = 0.46$	T1E	Bias*	$\gamma = 1$	$\gamma = 0.59$	$\gamma = 0.41$
correct	0.063	0.031	97.0	86.3	83.6	0.036	0.021	196.8	162.5	138.3
correct prior	0.071	0.021	96.5	85.7	82.2	0.038	0.017	196.5	161.9	136.5
correct strong prior	0.063	-0.014	97.3	86.6	83.8	0.046	-0.006	195.7	159.2	134.9
unadjusted	0.034	0.058	98.6	90.7	88.3	0.031	0.025	198.0	171.4	147.2
unadjusted prior	0.034	0.051	98.6	90.0	87.4	0.032	0.021	197.9	170.8	146.2
unadjusted strong prior	0.041	0.009	98.4	89.6	86.8	0.034	0.001	197.6	167.6	142.8
	Maximum sample size = 500					Maximum sample size = 1000				
			Expected sample size					Expected sample size		
Adjustment model	T1E	Bias*	$\gamma = 1$	$\gamma = 0.72$	$\gamma = 0.60$	T1E	Bias*	$\gamma = 1$	$\gamma = 0.80$	$\gamma = 0.72$
correct	0.028	0.010	493.7	404.9	334.9	0.022	0.008	992.0	823.2	681.5
correct prior	0.028	0.009	494.1	402.9	334.9	0.021	0.007	991.7	818.9	680.3
correct strong prior	0.028	0.004	493.7	401.2	329.9	0.023	0.005	990.5	813.9	675.7
unadjusted	0.026	0.016	494.6	426.0	367.0	0.024	0.010	990.2	859.6	734.2
unadjusted prior	0.029	0.015	493.4	426.7	364.9	0.024	0.009	989.5	858.6	733.7
unadjusted strong prior	0.031	0.009	492.6	423.0	362.7	0.023	0.007	989.5	858.1	733.9

A.6 Bias from overestimation in trials which stop early for superiority

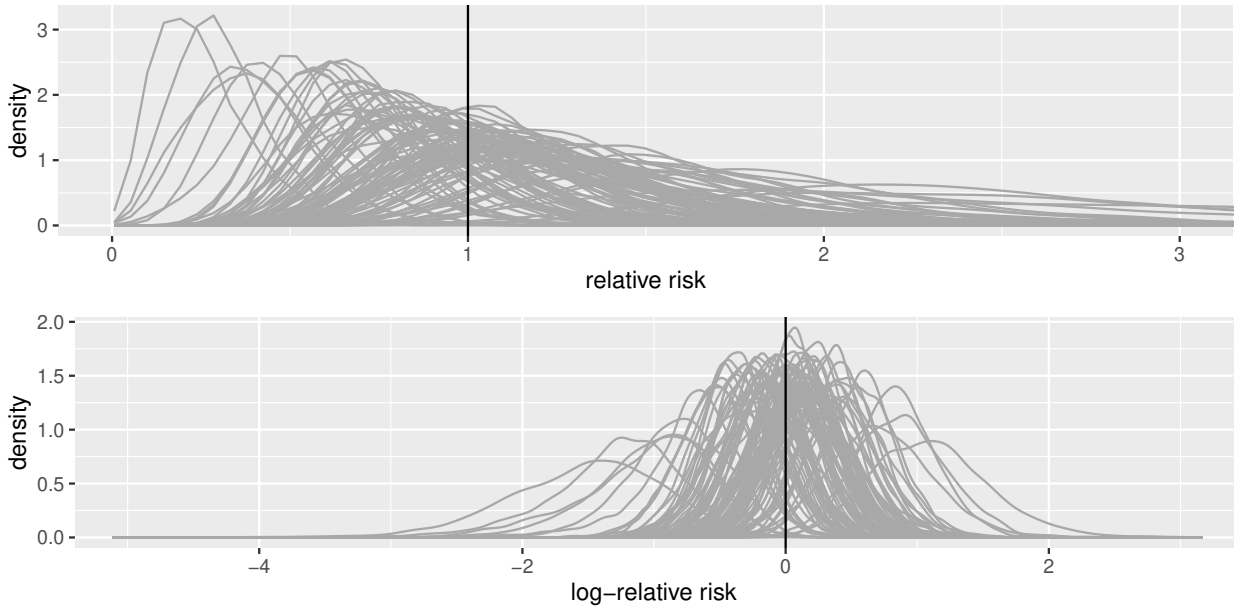

Figure A.6.1: 100 posterior distributions of the relative risk (top) and log-relative risk (bottom) for a trial with a binary endpoint and maximum sample size of 100. Vertical line represents the null treatment effect.

Walter et al. (2019) shows that overestimation is to be expected when trials permit early stopping for superiority. They consider the case of frequentist group-sequential designs and compare three different stopping rules which differ in how the overall α is divided among the interim and final analyses. The Pocock, O'Brien, and Fleming (PCK) stopping rule evenly divides α across all analyses (interim and final) keeping the stringency of the stopping criteria constant. This is the frequentist group sequential stopping rule most like the Bayesian stopping rule employed in the current manuscript, where a single value for the upper probability threshold u is used (u=0.99), thereby also keeping the stringency of the stopping rule constant across all interim and final analyses. It is shown that overestimation is expected for the PCK stopping rule, and so we conclude it should also be expected for the Bayesian stopping rule employed here, thus inducing the observed bias in the treatment

effect under the simulation scenarios. In Section 3.1 of Walter et al. (2019), the authors discuss observing greater over-estimation for the Haybittle and Peto (HP) stopping rule as compared to the PCK stopping rule. They state "rules (such as HP) that have a more stringent threshold for stopping involve a greater risk of over-estimation if the rule is actually invoked." This suggests that Bayesian stopping rules which have more stringent initial stopping criteria may lead to increased bias as compared to the stopping rules employed in the current manuscript, though we do not explore this further here.

Considering bias under the null, the difference in signs between the continuous endpoint versus the binary and time-to-event endpoints reflects the lower bounds of the estimands. The difference in means under the continuous endpoint is unbounded below, whereas the relative risk and hazard ratios are bounded below by zero. When bias is calculated on the log-relative risk and log-hazard ratio scales, most values for bias under the null for both the binary and time-to-event endpoints become negative as well, with greater bias for smaller sample sizes as in the continuous endpoint. This is explained visually in Figure A.6.1 and A.6.2, where 100 posterior distributions (Figure A.6.1) and posterior medians (Figure A.6.2) have been plotted for the null treatment effect for a trial with a binary endpoint under a maximum sample size of 100. The vertical lines represent the null values used for calculation of bias.

In the top panel of Figure A.6.1 on the relative risk scale, we see many right-skewed posteriors which lead to some posterior median estimates which are much greater than the null (Figure A.6.2, top panel). This induces positive bias under the null. When we move to the log scale, the right-skewed distributions become more normal in shape and are more centered around the null, but some of the posteriors which were closer to the lower boundary of 0 on the relative risk scale become left-skewed (Figure A.6.1, bottom panel). This results in some posterior medians becoming much less than the null (Figure A.6.2, bottom panel) which leads to negative bias under the null as in the continuous endpoint case. When calculating

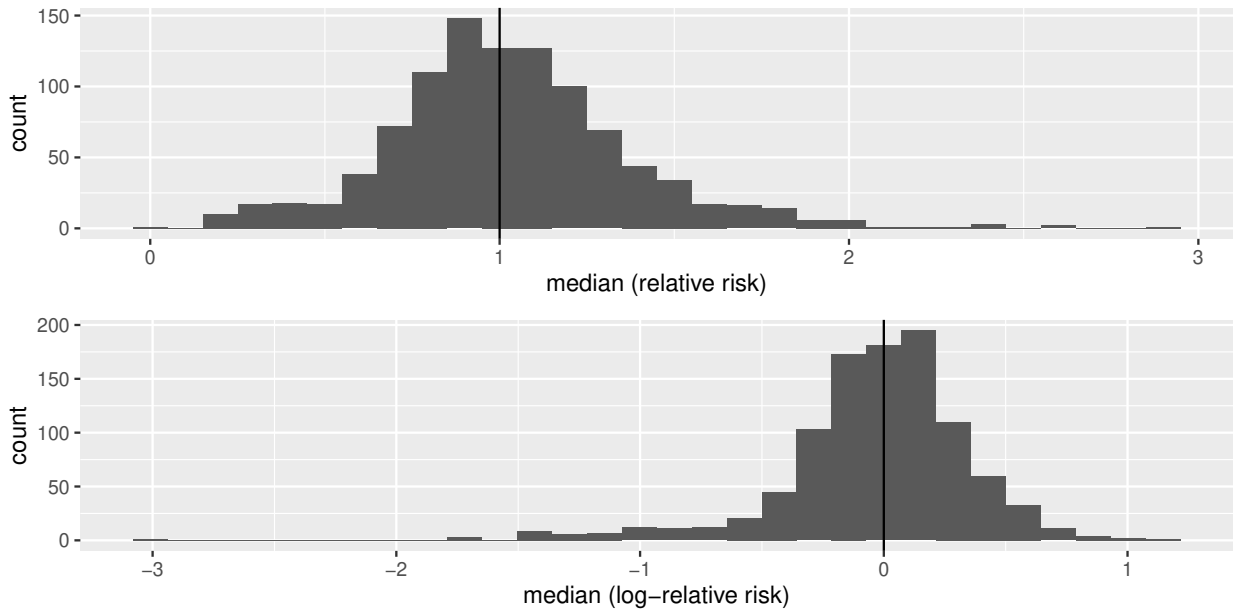

Figure A.6.2: 100 posterior medians of the relative risk (top) and log-relative risk (bottom) for a trial with a binary endpoint and maximum sample size of 100. Vertical line represents the null treatment effect.

bias for non-null treatment effects on the relative risk and hazard ratio scales, the posteriors are pushed further toward 0 than in the figures included in this section which is why they can still attain negative values and clearly exhibit overestimation in these cases.

A.7 Example of the non-collapsibility of the odds ratio

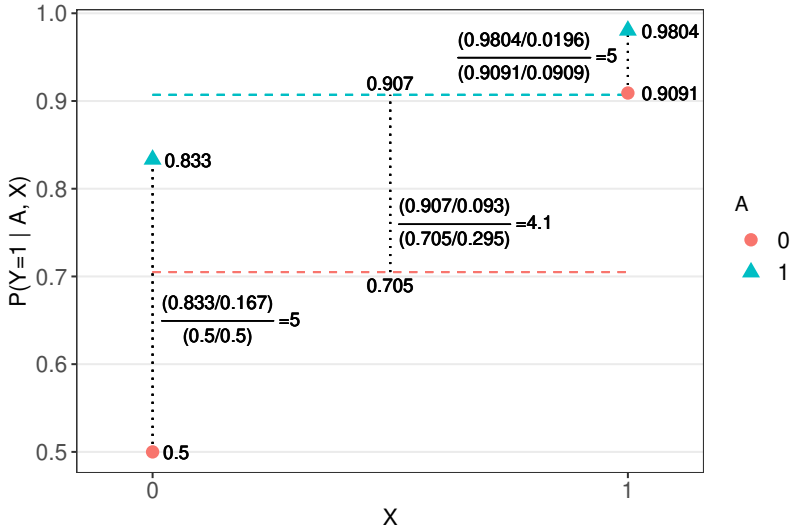

Figure A.7.1: Example of non-collapsibility of the odds ratio. Colored circles and triangles correspond to values of risk under treatment or control for different values of X.

We consider a slightly modified version of the example provided in Daniel et al. (2021). Consider a RCT with a binary endpoint following the logistic regression model below which contains a binary covariate X and binary treatment indicator A, and where $P(X = 1) = P(A = 1) = 0.5$. Define $\phi = \log(5)$, $\beta = \log(10)$, and $\boldsymbol{\theta} = \{\phi, \beta\}$. We assume the following model:

$$\text{logit}(P(Y = 1 \mid A, X)) = \phi A + \beta X.$$

In this example, the conditional odds ratio for those who are treated versus untreated is 5, regardless of the value of X. To see this, define $\mu(\boldsymbol{\theta}; A, X) = P(Y = 1 \mid A, X) = \text{logit}^{-1}(\phi A + \beta X)$. For $X = 0$, we have $\mu(\boldsymbol{\theta}; A = 1, X = 0) = 0.833$ and that $1 - \mu(\boldsymbol{\theta}; A = 1, X = 0) = 0.167$ yielding the odds for the event in those who are treated to be $\mu(\boldsymbol{\theta}; A = 1, X = 0)/1 - \mu(\boldsymbol{\theta}; A = 1, X = 0) = 0.833/0.167 = 5$. For those who are untreated, we have $\mu(\boldsymbol{\theta}; A = 0, X = 0) = 0.5$ and that $1 - \mu(\boldsymbol{\theta}; A = 0, X = 0) = 0.5$ yielding odds of $\mu(\boldsymbol{\theta}; A = 0, X = 0)/1 - \mu(\boldsymbol{\theta}; A = 0, X = 0) = 0.5/0.5 = 1$. Dividing these yields a conditional

odds ratio for those who are treated versus untreated under $X = 0$ to be $5/1 = 5$. Similar calculations for $X = 1$ yield $\mu(\boldsymbol{\theta}; A = 1, X = 1) = 0.9804$, $1 - \mu(\boldsymbol{\theta}; A = 1, X = 1) = 0.0196$ yielding the odds for the event in those who are treated to be $0.9804/0.0196 = 50$. For those who are untreated, we have $\mu(\boldsymbol{\theta}; A = 0, X = 1) = 0.9091$ and that $1 - \mu(\boldsymbol{\theta}; A = 0, X = 1) = 0.0909$ yielding odds of $0.9091/0.0909 = 10$. Dividing these yields a conditional odds ratio for those who are treated versus untreated under $X = 1$ to be $50/10 = 5$. In Figure A.7.1, we that these conditional odds ratios correspond to a vertical comparison of the risks under the treatment assignments A for either value of X (dotted vertical lines). To obtain the marginal odds ratio, we must average these risks with respect to the distribution of X. This yields the horizontal dashed lines (colored by value of A) where $\mu(\boldsymbol{\theta}; A = 1) = 0.5(0.833) + 0.5(0.980) = 0.907$ and $\mu(\boldsymbol{\theta}; A = 0) = 0.5(0.5) + 0.5(0.909) = 0.705$. The marginal odds ratio then corresponds to a vertical comparison of these horizontal lines. Doing so yields a marginal odds ratio of $\gamma(\boldsymbol{\theta}) = (0.907/0.093)/(0.705/0.295) = 4.1$. We see that the marginal odds ratio is not equal to the conditional odds ratio, and thus the odds ratio is non-collapsible.

This same example can be viewed using two-by-two tables, where the cells contain the proportions expected under each combination of treatment assignment and covariate value.

X = 0

	$A = 1$	$A = 0$	$P(Y = y)$
$Y = 1$	0.8333333	0.5	0.667
$Y = 0$	0.16666667	0.5	0.333
$P(A = a)$	0.5	0.5	1

X = 1

	$A = 1$	$A = 0$	$P(Y = y)$
$Y = 1$	0.9803922	0.9090909	0.945
$Y = 0$	0.01960784	0.09090909	0.055
$P(A = a)$	0.5	0.5	1

The conditional odds ratio from each table is 5. Below, we consider the marginal table with proportions expected under each combination of treatment assignment and covariate value. This is found by averaging the risk values in the conditional tables with respect to

the distribution of X, where we recall that $P(X = 1) = 0.5$.

	$A = 1$	$A = 0$	$P(Y = y)$
$Y = 1$	0.90686275	0.70454545	0.806
$Y = 0$	0.09313725	0.29545455	0.194
$P(A = a)$	0.5	0.5	1

The marginal odds ratio from the table above is 4.1. We observe that the true conditional and marginal odds ratios are not equal, thus showing the odds ratio is non-collapsible.

References

Aishah, A., Loffler, K. A., Toson, B., Mukherjee, S., Adams, R. J., Altree, T. J., Ainge-Allen, H. W., Yee, B. J., Grunstein, R. R., Carberry, J. C., et al. (2023). One Month Dosing of Atomoxetine Plus Oxybutynin in Obstructive Sleep Apnea: A Randomized, Placebo-Controlled Trial. *Annals of the American Thoracic Society*, 20(4):584–595.

Anscombe, F. (1963). Sequential Medical Trials. *Journal of the American Statistical Association*, 58(302):365–383.

Aranda-Ordaz, F. J. (1981). On Two Families of Transformations to Additivity for Binary Response Data. *Biometrika*, 68(2):357–363.

Austin, P. C., Manca, A., Zwarenstein, M., Juurlink, D. N., and Stanbrook, M. B. (2010). A Substantial and Confusing Variation Exists in Handling of Baseline Covariates in Randomized Controlled Trials: A Review of Trials Published in Leading Medical Journals. *Journal of Clinical Epidemiology*, 63(2):142–153.

Avorn, J. (2015). The $2.6 Billion Pill — Methodologic and Policy Considerations. *New England Journal of Medicine*, 372(20):1877–1879.

Babb, J., Rogatko, A., and Zacks, S. (1998). Cancer Phase I Clinical Trials: Efficient Dose Escalation with Overdose Control. *Statistics in Medicine*, 17(10):1103–1120.

Babb, J. S. and Rogatko, A. (2001). Patient Specific Dosing in a Cancer Phase I Clinical Trial. *Statistics in Medicine*, 20(14):2079–2090.

Barker, A., Sigman, C., Kelloff, G. J., Hylton, N., Berry, D. A., and Esserman, L. (2009). I-SPY 2: an Adaptive Breast Cancer Trial Design in the Setting of Neoadjuvant Chemotherapy. *Clinical Pharmacology & Therapeutics*, 86(1):97–100.

Benjafield, A. V., Ayas, N. T., Eastwood, P. R., Heinzer, R., Ip, M. S., Morrell, M. J., Nunez, C. M., Patel, S. R., Penzel, T., Pépin, J.-L., et al. (2019). Estimation of the Global Prevalence and Burden of Obstructive Sleep Apnoea: A Literature-Based Analysis. *The Lancet Respiratory Medicine*, 7(8):687–698.

Benkeser, D., Díaz, I., Luedtke, A., Segal, J., Scharfstein, D., and Rosenblum, M. (2021). Improving Precision and Power in Randomized Trials for COVID-19 Treatments Using Covariate Adjustment, for Binary, Ordinal, and Time-to-event Outcomes. *Biometrics*, 77(4):1467–1481.

Berger, J. O. and Wolpert, R. L. (1988). *The Likelihood Principle*. Institute of Mathematical Statistics, Hayward, California.

Bernardo, J. M. and Smith, A. F. (1994). *Bayesian Theory*. John Wiley & Sons, Chichester, England.

Berry, D. A. (1985). Interim Analyses in Clinical Trials: Classical vs. Bayesian Approaches. *Statistics in Medicine*, 4(4):521–526.

Berry, D. A. (1987). Interim Analysis in Clinical Trials: the Role of the Likelihood Principle. *The American Statistician*, 41(2):117–122.

Berry, D. A. and Eick, S. G. (1995). Adaptive Assignment versus Balanced Randomization in Clinical Trials: A Decision Analysis. *Statistics in Medicine*, 14(3):231–246.

Berry, D. A. and Ho, C.-H. (1988). One-Sided Sequential Stopping Boundaries for Clinical Trials: A Decision-Theoretic Approach. *Biometrics*, 44(1):219–227.

Berry, S. M., Carlin, B. P., Lee, J. J., and Muller, P. (2010). *Bayesian Adaptive Methods for Clinical Trials*. CRC Press, Boca Raton, Florida.

Binois, M. and Gramacy, R. B. (2021). hetgp: Heteroskedastic Gaussian Process Modeling and Sequential Design in R. *Journal of Statistical Software*, 98(13):1–44.

Binois, M., Gramacy, R. B., and Ludkovski, M. (2018). Practical Heteroscedastic Gaussian Process Modeling for Large Simulation Experiments. *Journal of Computational and Graphical Statistics*, 27(4):808–821.

Bradford Hill, A. (1962). *Statistical Methods in Clinical and Preventive Medicine*. E&S Livingstone, Edinburgh, Scotland.

Brilleman, S. L., Elci, E. M., Novik, J. B., and Wolfe, R. (2020). Bayesian Survival Analysis Using the rstanarm R Package. *arXiv preprint arXiv:2002.09633*.

Bull, A. D. (2011). Convergence Rates of Efficient Global Optimization Algorithms. *Journal of Machine Learning Research*, 12(88):2879–2904.

Bull, J. P. (1959). The Historical Development of Clinical Therapeutic Trials. *Journal of Chronic Diseases*, 10(3):218–248.

Byar, D. P. and Corle, D. K. (1977). Selecting Optimal Treatment in Clinical Trials using Covariate Information. *Journal of Chronic Diseases*, 30(7):445–459.

Cai, C., Yuan, Y., and Ji, Y. (2014). A Bayesian Dose-Finding Design for Oncology Clinical Trials of Combinational Biological Agents. *Journal of the Royal Statistical Society, Series C*, 63(1):159–173.

Carlin, B. P., Kadane, J. B., and Gelfand, A. E. (1998). Approaches for Optimal Sequential Decision Analysis in Clinical Trials. *Biometrics*, 54(3):964–975.

Chaloner, K. and Verdinelli, I. (1995). Bayesian Experimental Design: A Review. *Statistical Science*, 10(3):273–304.

Chapple, A. G. and Thall, P. F. (2019). A Hybrid Phase I-II/III Clinical Trial Design Allowing Dose Re-Optimization in Phase III. *Biometrics*, 75(2):371–381.

Chen, Z., Krailo, M. D., Azen, S. P., and Tighiouart, M. (2010). A Novel Toxicity Scoring System Treating Toxicity Response as a Quasi-Continuous Variable in Phase I Clinical Trials. *Contemporary Clinical Trials*, 31(5):473–482.

Chen, Z., Tighiouart, M., and Kowalski, J. (2012). Dose Escalation with Overdose Control Using a Quasi-Continuous Toxicity Score in Cancer Phase I Clinical Trials. *Contemporary Clinical Trials*, 33(5):949–958.

Chipman, H. A., George, E. I., and McCulloch, R. E. (2010). BART: Bayesian Additive Regression Trees. *The Annals of Applied Statistics*, 4(1):266–298.

Ciolino, J. D., Palac, H. L., Yang, A., Vaca, M., and Belli, H. M. (2019). Ideal vs. Real: A Systematic Review on Handling Covariates in Randomized Controlled Trials. *BMC Medical Research Methodology*, 19(1):1–11.

Committee for Proprietary Medicinal Products(2004). Points to Consider on Adjustment for Baseline Covariates. *Statistics in Medicine*, 23(5):701–709.

Cornfield, J. (1966). Sequential Trials, Sequential Analysis and the Likelihood Principle. *The American Statistician*, 20(2):18–23.

Cox, D. R. (1957). The use of a Concomitant Variable in Selecting an Experimental Design. *Biometrika*, 44(1/2):150–158.

Cox, D. R. and McCullagh, P. (1982). Some Aspects of Analysis of Covariance. *Biometrics*, 38(3):541–561.

Cunanan, K. and Koopmeiners, J. S. (2014). Evaluating the Performance of Copula Models in Phase I-II Clinical Trials under Model Misspecification. *BMC Medical Research Methodology*, 14:1–11.

Daniel, R., Zhang, J., and Farewell, D. (2021). Making Apples from Oranges: Comparing Noncollapsible Effect Estimators and their Standard Errors after Adjustment for Different Covariate Sets. *Biometrical Journal*, 63(3):528–557.

Daniels, M. J., Linero, A., and Roy, J. (2023). *Bayesian Nonparametrics for Causal Inference and Missing Data*. CRC Press, Boca Raton, Florida.

Dixon, D. O. and Simon, R. (1991). Bayesian Subset Analysis. *Biometrics*, 47(3):871–881.

Duvenaud, D. K., Nickisch, H., and Rasmussen, C. (2011). Additive Gaussian Processes. In *The 25th Annual Neural Information Processing Systems Conference*, Granada, Spain.

Eichhorn, B. H. and Zacks, S. (1973). Sequential Search of an Optimal Dosage, I. *Journal of the American Statistical Association*, 68(343):594–598.

European Medicines Agency (2015). Guideline on Adjustment for Baseline Covariates in Clinical Trials. https://www.ema.europa.eu/en/documents/scientific-guideline/guideline-adjustment-baseline-covariates-clinical-trials_en.pdf (accessed 12 May 2022).

Farewell, V. T., Tom, B. D. M., and Royston, P. (2004). The Impact of Dichotomization on the Efficiency of Testing for an Interaction Effect in Exponential Family Models. *Journal of the American Statistical Association*, 99(467):822–831.

Fisher, R. A. (1925). *Statistical Methods for Research Workers*. Oliver and Boyd, Edinburgh, Scotland.

Fisher, R. A. and Mackenzie, W. A. (1923). Studies in Crop Variation. II. The Manurial Response of Different Potato Varieties. *The Journal of Agricultural Science*, 13(3):311–320.

Ford, I., Norrie, J., and Ahmadi, S. (1995). Model Inconsistency, Illustrated by the Cox Proportional Hazards Model. *Statistics in Medicine*, 14(8):735–746.

Freeman, N. L., Browder, S. E., McGinigle, K. L., and Kosorok, M. R. (2022). Dynamic Treatment Regime Characterization via Value Function Surrogate with an Application to Partial Compliance. *arXiv preprint arXiv:2212.00650*.

Freeman Jr, D. H. and Holford, T. R. (1980). Summary Rates. *Biometrics*, 36(2):195–205.

Friedman, L. M., Furberg, C. D., DeMets, D. L., Reboussin, D. M., and Granger, C. B. (2015). *Fundamentals of Clinical Trials*. Springer, Cham, Switzerland.

Gabry, J. and Goodrich, B. (2022). Prior Distributions for rstanarm Models. https://cran.r-project.org/web/packages/rstanarm/vignettes/priors.html (accessed 21 April 2022).

Gail, M. and Simon, R. (1985). Testing for Qualitative Interactions between Treatment Effects and Patient Subsets. *Biometrics*, 41(2):361–372.

Gardner, J. R., Kusner, M. J., Xu, Z. E., Weinberger, K. Q., and Cunningham, J. P. (2014). Bayesian Optimization with Inequality Constraints. In *Proceedings of the 31st International Conference on Machine Learning*, Beijing, China. 32:937–945.

Garnett, R. (2023). *Bayesian Optimization*. Cambridge University Press, Cambridge, UK.

Gelbart, M. A., Snoek, J., and Adams, R. P. (2014). Bayesian Optimization with Unknown Constraints. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, Quebec City, Canada. 30:250–259.

Gelfand, A. E. and Smith, A. F. (1990). Sampling-Based Approaches to Calculating Marginal Densities. *Journal of the American Statistical Association*, 85(410):398–409.

Gelman, A., Bois, F., and Jiang, J. (1996). Physiological Pharmacokinetic Analysis using Population Modeling and Informative Prior Distributions. *Journal of the American Statistical Association*, 91(436):1400–1412.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2014). *Bayesian Data Analysis*. CRC Press, Boca Raton, Florida.

Gelman, A. and Rubin, D. B. (1992). Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4):457–472.

Giovagnoli, A. (2021). The Bayesian Design of Adaptive Clinical Trials. *International Journal of Environmental Research and Public Health*, 18(2):530.

Golchi, S., Bingham, D. R., Chipman, H., and Campbell, D. A. (2015). Monotone Emulation of Computer Experiments. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):370–392.

Golchi, S. and Thorlund, K. (2020). Sequential Monte Carlo for Response Adaptive Randomized Trials. *Biostatistics*, 21(2):287–301.

Goodrich, B., Gabry, J., Ali, I., and Brilleman, S. (2020). rstanarm: Bayesian Applied Regression Modeling via Stan. https://mc-stan.org/rstanarm (version 2.21.2).

Gramacy, R. B. (2020). *Surrogates: Gaussian Process Modeling, Design and Optimization for the Applied Sciences*. Chapman Hall/CRC, Boca Raton, Florida.

Gramacy, R. B., Gray, G. A., Le Digabel, S., Lee, H. K., Ranjan, P., Wells, G., and Wild, S. M. (2016). Modeling an Augmented Lagrangian for Blackbox Constrained Optimization. *Technometrics*, 58(1):1–11.

Gramacy, R. B. and Lee, H. K. H. (2011). Optimization Under Unknown Constraints. In *Bayesian Statistics*. Oxford University Press, Oxford, UK. 9:229–256.

Gramacy, R. B. and Polson, N. G. (2011). Particle Learning of Gaussian Process Models for Sequential Design and Optimization. *Journal of Computational and Graphical Statistics*, 20(1):102–118.

Guo, B. and Zang, Y. (2022). A Bayesian Phase I/II Biomarker-based Design for Identifying Subgroup-Specific Optimal Dose for Immunotherapy. *Statistical Methods in Medical Research*, 31(6):1104–1119.

Harrell, F. and Slaughter, J. (2021). Biostatistics for Biomedical Research. http://hbiostat.org/doc/bbr.pdf (accessed 03 Feb 2022).

Hauck, W. W., Anderson, S., and Marcus, S. M. (1998). Should We Adjust for Covariates in Nonlinear Regression Analyses of Randomized Trials? *Controlled Clinical Trials*, 19(3):249–256.

Hernández, A. V., Eijkemans, M. J., and Steyerberg, E. W. (2006). Randomized Controlled Trials With Time-to-Event Outcomes: How Much Does Prespecified Covariate Adjustment Increase Power? *Annals of Epidemiology*, 16(1):41–48.

Hernández, A. V., Steyerberg, E. W., and Habbema, J. F. (2004). Covariate Adjustment in Randomized Controlled Trials with Dichotomous Outcomes Increases Statistical Power and Reduces Sample Size Requirements. *Journal of Clinical Epidemiology*, 57(5):454–460.

Herson, J. (1979). Predictive Probability Early Termination Plans for Phase II Clinical Trials. *Biometrics*, 35(4):775–783.

Hirakawa, A., Wages, N. A., Sato, H., and Matsui, S. (2015). A Comparative Study of Adaptive Dose-Finding Designs for Phase I Oncology Trials of Combination Therapies. *Statistics in Medicine*, 34(24):3194–3213.

Hohl, C. M., Rosychuk, R. J., Archambault, P. M., O'Sullivan, F., Leeies, M., Mercier, É., Clark, G., Innes, G. D., Brooks, S. C., Hayward, J., Ho, V., Jelic, T., Welsford, M., Sivilotti, M. L. A., Morrison, L. J., and Perry, J. J. (2022a). The CCEDRRN COVID-19 Mortality Score to Predict Death Among Nonpalliative Patients with COVID-19 Presenting to Emergency Departments: A Derivation and Validation Study. *Canadian Medical Association Open Access Journal*, 10(1):E90–E99.

Hohl, C. M., Rosychuk, R. J., Hau, J. P., Hayward, J., Landes, M., Yan, J. W., Ting, D. K., Welsford, M., Archambault, P. M., Mercier, E., Chandra, K., Davis, P., Vaillancourt, S., Leeies, M., Small, S., and Morrison, L. J. (2022b). Treatments, Resource

Utilization, and Outcomes of COVID-19 Patients Presenting to Emergency Departments Across Pandemic Waves: an Observational Study by the Canadian COVID-19 Emergency Department Rapid Response Network (CCEDRRN). *Canadian Journal of Emergency Medicine*, 24:397–407.

Hohl, C. M., Rosychuk, R. J., McRae, A. D., Brooks, S. C., Archambault, P., Fok, P. T., Davis, P., Jelic, T., Turner, J. P., Rowe, B. H., Mercier, É., Cheng, I., Taylor, J., Daoust, R., Ohle, R., Andolfatto, G., Atzema, C., Hayward, J., Khangura, J. K., Landes, M., Lang, E., Martin, I., Mohindra, R., Ting, D. K., Vaillancourt, S., Welsford, M., Brar, B., Dahn, T., Wiemer, H., Yadav, K., Yan, J. W., Stachura, M., McGavin, C., Perry, J. J., and Morrison, L. J. (2021). Development of the Canadian COVID-19 Emergency Department Rapid Response Network Population-based Registry: A Methodology Study. *Canadian Medical Association Open Access Journal*, 9(1):E261–E270.

Houede, N., Thall, P. F., Nguyen, H., Paoletti, X., and Kramar, A. (2010). Utility-Based Optimization of Combination Therapy Using Ordinal Toxicity and Efficacy in Phase I/II Trials. *Biometrics*, 66(2):532–540.

Huang, D., Allen, T. T., Notz, W. I., and Zeng, N. (2006). Global Optimization of Stochastic Black-Box Systems via Sequential Kriging Meta-Models. *Journal of Global Optimization*, 34:441–466.

Ibrahim, J. G. and Chen, M.-H. (2000). Power Prior Distributions for Regression Models. *Statistical Science*, 15(1):46–60.

Jones, D. R., Schonlau, M., and Welch, W. J. (1998). Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13(4):455–492.

Kahan, B. C., Jairath, V., Doré, C. J., and Morris, T. P. (2014). The Risks and Rewards of Covariate Adjustment in Randomized Trials: an Assessment of 12 Outcomes from 8 Studies. *Trials*, 15(1):1–7.

Kalton, G. (1968). Standardization: A Technique to Control for Extraneous Variables. *Journal of the Royal Statistical Society, Series C*, 17(2):118–136.

Keil, A. P., Daza, E. J., Engel, S. M., Buckley, J. P., and Edwards, J. K. (2018). A Bayesian Approach to the G-formula. *Statistical Methods in Medical Research*, 27(10):3183–3204.

Kimko, H. and Pinheiro, J. (2015). Model-Based Clinical Drug Development in the Past, Present and Future: A Commentary. *British Journal of Clinical Pharmacology*, 79(1):108–116.

Korn, E. L. (2004). Nontoxicity Endpoints in Phase I Trial Designs for Targeted, Non-Cytotoxic Agents. *Journal of the National Cancer Institute*, 96(13):977–978.

Kushner, H. (1964). A New Method of Locating the Maximum Point of an Arbitrary Multipeak Curve in the Presence of Noise. *Journal of Basic Engineering*, 86(1):97–106.

Lane, P. W. and Nelder, J. A. (1982). Analysis of Covariance and Standardization as Instances of Prediction. *Biometrics*, 38(3):613–621.

Le-Rademacher, J. G., Hillman, S., Storrick, E., Mahoney, M. R., Thall, P. F., Jatoi, A., and Mandrekar, S. J. (2020). Adverse Event Burden Score — a Versatile Summary Measure for Cancer Clinical Trials. *Cancers*, 12(11):3251.

Le Tourneau, C., Lee, J. J., and Siu, L. L. (2009). Dose Escalation Methods in Phase I Cancer Clinical Trials. *JNCI: Journal of the National Cancer Institute*, 101(10):708–720.

Lee, K. M., Robertson, D. S., Jaki, T., and Emsley, R. (2022). The Benefits of Covariate Adjustment for Adaptive Multi-Arm Designs. *Statistical Methods in Medical Research*, 31(11):2104–2121.

Lee, S., Hershman, D., Martin, P., Leonard, J., and Cheung, Y. (2012). Toxicity Burden Score: A Novel Approach to Summarize Multiple Toxic Effects. *Annals of Oncology*, 23(2):537–541.

Leske, M. C., Heijl, A., Hussein, M., Bengtsson, B., Hyman, L., Komaroff, E., Group, E. M. G. T., et al. (2003). Factors for Glaucoma Progression and the Effect of Treatment: the Early Manifest Glaucoma Trial. *Archives of Ophthalmology*, 121(1):48–56.

Letham, B., Karrer, B., Ottoni, G., and Bakshy, E. (2019). Constrained Bayesian Optimization with Noisy Experiments. *Bayesian Analysis*, 14(2):495–519.

Lewis, J. A. (1999). Statistical Principles for Clinical Trials (ICH E9): An Introductory Note on an International Guideline. *Statistics in Medicine*, 18(15):1903–1942.

Li, D. H., Whitmore, J. B., Guo, W., and Ji, Y. (2017). Toxicity and Efficacy Probability Interval Design for Phase I Adoptive Cell Therapy Dose-Finding Clinical Trials. *Clinical Cancer Research*, 23(1):13–20.

Lin, L. and Dunson, D. B. (2014). Bayesian Monotone Regression Using Gaussian Process Projection. *Biometrika*, 101(2):303–317.

Lind, J. (1753). *A Treatise of the Scurvy*. Sands, Murray, and Cochran, Edinburgh, Scotland.

Lindsey, J. K. and Lambert, P. (1998). On the Appropriateness of Marginal Models for Repeated Measurements in Clinical Trials. *Statistics in Medicine*, 17(4):447–469.

Linero, A. R. and Antonelli, J. L. (2023). The How and Why of Bayesian Nonparametric Causal Inference. *Wiley Interdisciplinary Reviews: Computational Statistics*, 15(1):e1583.

Lyu, J., Ji, Y., Zhao, N., and Catenacci, D. V. (2019). AAA: Triple Adaptive Bayesian Designs for the Identification of Optimal Dose Combinations in Dual-Agent Dose Finding Trials. *Journal of the Royal Statistical Society, Series C*, 68(2):385–410.

Medical Research Council (1948). Streptomycin Treatment of Pulmonary Tuberculosis: A Medical Research Council Investigation. *British Medical Journal*, 2:769–782.

Morgan-Wall, T. (2022). *spacefillr: Space-Filling Random and Quasi-Random Sequences*. R package version 0.3.2. https://CRAN.R-project.org/package=spacefillr.

Morita, S., Thall, P. F., and Takeda, K. (2017). A Simulation Study of Methods for Selecting Subgroup-Specific Doses in Phase 1 Trials. *Pharmaceutical Statistics*, 16(2):143–156.

Morris, T. P., Walker, A. S., Williamson, E. J., and White, I. R. (2022). Planning a Method for Covariate Adjustment in Individually Randomised Trials: A Practical Guide. *Trials*, 23(1):328.

Mozgunov, P., Cro, S., Lingford-Hughes, A., Paterson, L. M., and Jaki, T. (2022). A Dose-Finding Design for Dual-Agent Trials with Patient-Specific Doses for One Agent with Application to an Opiate Detoxification Trial. *Pharmaceutical Statistics*, 21(2):476–495.

Mozgunov, P. and Jaki, T. (2019a). A Flexible Design for Advanced Phase I/II Clinical Trials with Continuous Efficacy Endpoints. *Biometrical Journal*, 61(6):1477–1492.

Mozgunov, P. and Jaki, T. (2019b). An Information Theoretic Phase I–II Design for Molecularly Targeted Agents that does not Require an Assumption of Monotonicity. *Journal of the Royal Statistical Society, Series C*, 68(2):347–367.

Mueller, P. S., Montori, V. M., Bassler, D., Koenig, B. A., and Guyatt, G. H. (2007). Ethical Issues in Stopping Randomized Trials Early Because of Apparent Benefit. *Annals of Internal Medicine*, 146(12):878–881.

Murphy, K. P. (2023). *Probabilistic Machine Learning: Advanced Topics*. MIT Press, Cambridge, Massachusetts. http://probml.github.io/book2.

Oganisian, A. and Roy, J. A. (2021). A Practical Introduction to Bayesian Estimation of Causal Effects: Parametric and Nonparametric Approaches. *Statistics in Medicine*, 40(2):518–551.

O'Quigley, J., Pepe, M., and Fisher, L. (1990). Continual Reassessment Method: A Practical Design for Phase 1 Clinical Trials in Cancer. *Biometrics*, 46(1):33–48.

Picheny, V., Ginsbourger, D., Richet, Y., and Caplin, G. (2013a). Quantile-Based Optimization of Noisy Computer Experiments with Tunable Precision. *Technometrics*, 55(1):2–13.

Picheny, V., Wagner, T., and Ginsbourger, D. (2013b). A Benchmark of Kriging-Based Infill Criteria for Noisy Optimization. *Structural and Multidisciplinary Optimization*, 48:607–626.

Pinheiro, J. C., Bretz, F., and Hsu, C.-H. (2011). Adaptive Trial Designs. In *Clinical Trial Simulations: Applications and Trends*, New York, New York. Springer. 1:109–130.

Pocock, S. J. (1983). *Clinical Trials: A Practical Approach*. John Wiley & Sons, Chichester, England.

Pocock, S. J., Assmann, S. E., Enos, L. E., and Kasten, L. E. (2002). Subgroup Analysis, Covariate Adjustment and Baseline Comparisons in Clinical Trial Reporting: Current Practice and Problems. *Statistics in Medicine*, 21(19):2917–2930.

Psioda, M. A., Xu, J., Jiang, Q., Ke, C., Yang, Z., and Ibrahim, J. G. (2021). Bayesian Adaptive Basket Trial Design using Model Averaging. *Biostatistics*, 22(1):19–34.

R Core Team (2022). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/.

Remiro-Azócar, A., Heath, A., and Baio, G. (2022). Parametric G-computation for Compatible Indirect Treatment Comparisons with Limited Individual Patient Data. *Research Synthesis Methods*, 13(6):716–744.

Robertson, D. S., Choodari-Oskooei, B., Dimairo, M., Flight, L., Pallmann, P., and Jaki, T. (2023). Point Estimation for Adaptive Trial Designs I: A Methodological Review. *Statistics in Medicine*, 42(2):122–145.

Robinson, L. D. and Jewell, N. P. (1991). Some Surprising Results about Covariate Adjustment in Logistic Regression Models. *International Statistical Review*, 59(2):227–240.

Rodriguez Duque, D., Stephens, D. A., and Moodie, E. E.M. (2022). Estimation of Optimal Dynamic Treatment Regimes using Gaussian Process Emulation. *arXiv preprint arXiv:2105.12259*.

Rosenberg, R., Abaluck, B., and Thein, S. (2022). Combination of Atomoxetine with the Novel Antimuscarinic Aroxybutynin Improves Mild to Moderate OSA. *Journal of Clinical Sleep Medicine*, 18(12):2837–2844.

Rosenberger, W. F. and Sverdlov, O. (2008). Handling Covariates in the Design of Clinical Trials. *Statistical Science*, 23(3):404–419.

Rotenberg, B. W., Murariu, D., and Pang, K. P. (2016). Trends in CPAP Adherence Over Twenty Years of Data Collection: A Flattened Curve. *Journal of Otolaryngology-Head & Neck Surgery*, 45(1):43.

Royston, P., Altman, D. G., and Sauerbrei, W. (2006). Dichotomizing Continuous Predictors in Multiple Regression: A Bad Idea. *Statistics in Medicine*, 25(1):127–141.

Rubin, D. B. (1981). The Bayesian Bootstrap. *The Annals of Statistics*, 9(1):130–134.

Russek-Cohen, E. and Simon, R. M. (1997). Evaluating Treatments when a Gender by Treatment Interaction may Exist. *Statistics in Medicine*, 16(4):455–464.

Saarela, O., Arjas, E., Stephens, D. A., and Moodie, E. (2015a). Predictive Bayesian Inference and Dynamic Treatment Regimes. *Biometrical Journal*, 57(6):941–958.

Saarela, O., Stephens, D. A., Moodie, E. E., and Klein, M. B. (2015b). On Bayesian Estimation of Marginal Structural Models. *Biometrics*, 71(2):279–288.

Sauer, A., Gramacy, R. B., and Higdon, D. (2023). Active Learning for Deep Gaussian Process Surrogates. *Technometrics*, 65(1):4–18.

Schweitzer, P. K., Taranto-Montemurro, L., Ojile, J. M., Thein, S. G., Drake, C. L., Rosenberg, R., Corser, B., Abaluck, B., Sangal, R. B., and Maynard, J. (2023). The Combination

of Aroxybutynin and Atomoxetine in the Treatment of Obstructive Sleep Apnea (MARI-POSA): A Randomized Controlled Trial. *American Journal of Respiratory and Critical Care Medicine*, 208(12):1316–1327.

Senn, S. (1989). Covariate Imbalance and Random Allocation in Clinical Trials. *Statistics in Medicine*, 8(4):467–475.

Senn, S. (2013). Seven Myths of Randomisation in Clinical Trials. *Statistics in Medicine*, 32(9):1439–1450.

Simon, R. (1982). Patient Subsets and Variation in Therapeutic Efficacy. *British Journal of Clinical Pharmacology*, 14(4):473–482.

Simon, R. (2002). Bayesian Subset Analysis: Application to Studying Treatment-by-Gender Interactions. *Statistics in Medicine*, 21(19):2909–2916.

Spiegelhalter, D. J., Abrams, K. R., and Myles, J. P. (2004). *Bayesian Approaches to Clinical Trials and Health-Care Evaluation*. John Wiley & Sons, Chichester, England.

Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. (2010). Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. In *Proceedings of the 27th International Conference on Machine Learning*, Haifa, Israel. 28:1015–1022.

Stephens, D., Nobre, W., Moodie, E., and Schmidt, A. (2023). Causal Inference Under Misspecification: Adjustment Based on the Propensity Score. *Bayesian Analysis*, 18(2):639–694.

Stitelman, O. M., Wester, C. W., De Gruttola, V., and van der Laan, M. J. (2011). Targeted Maximum Likelihood Estimation of Effect Modification Parameters in Survival Analysis. *The International Journal of Biostatistics*, 7(1):19.

Storer, B. E. (1989). Design and Analysis of Phase I Clinical Trials. *Biometrics*, 45(3):925–937.

Sweeting, M. J. and Mander, A. P. (2012). Escalation Strategies for Combination Therapy Phase I Trials. *Pharmaceutical Statistics*, 11(3):258–266.

Takahashi, A. and Suzuki, T. (2021a). Bayesian Optimization Design for Dose-Finding Based on Toxicity and Efficacy Outcomes in Phase I/II Clinical Trials. *Pharmaceutical Statistics*, 20(3):422–439.

Takahashi, A. and Suzuki, T. (2021b). Bayesian Optimization for Estimating the Maximum Tolerated Dose in Phase I Clinical Trials. *Contemporary Clinical Trials Communications*, 21:100753.

Thall, P. F. (2008). A Review of Phase 2–3 Clinical Trial Designs. *Lifetime Data Analysis*, 14(1):37–53.

Thall, P. F. (2021). Adaptive Enrichment Designs in Clinical Trials. *Annual Review of Statistics and its Application*, 8(1):393–411.

Thall, P. F. and Cook, J. D. (2004). Dose-Finding Based on Efficacy-Toxicity Trade-Offs. *Biometrics*, 60(3):684–693.

Thall, P. F. and Simon, R. (1994). A Bayesian Approach to Establishang Sample Size and Monitoring Criteria for Phase II Clinical Trials. *Controlled Clinical Trials*, 15(6):463–481.

Thall, P. F., Simon, R. M., and Estey, E. H. (1995). Bayesian Sequential Monitoring Designs for Single-Arm Clinical Trials with Multiple Outcomes. *Statistics in Medicine*, 14(4):357–379.

Thall, P. F. and Wathen, J. K. (2007). Practical Bayesian Adaptive Randomisation in Clinical Trials. *European Journal of Cancer*, 43(5):859–866.

Thall, P. F., Wooten, L. H., and Tannir, N. M. (2005). Monitoring Event Times in Early Phase Clinical Trials: Some Practical Issues. *Clinical Trials*, 2(6):467–478.

Thatcher, A. R. (1964). Relationships Between Bayesian and Confidence Limits for Predictions. *Journal of the Royal Statistical Society, Series B*, 26(2):176–192.

Thomas, N., Sweeney, K., and Somayaji, V. (2014). Meta-Analysis of Clinical Dose-Response in a Large Drug Development Portfolio. *Statistics in Biopharmaceutical Research*, 6(4):302–317.

Tracey, B. D. and Wolpert, D. (2018). Upgrading from Gaussian Processes to Student's-T Processes. In *2018 AIAA Non-Deterministic Approaches Conference*, Kissimmee, Florida.

Uemura, K., Ando, Y., and Matsuyama, Y. (2017). Utility of Adaptive Sample Size Designs and a Review Example. *Journal of Statistical Science and Application*, 5(1-2):1–15.

United States Food and Drug Administration (2018a). Precision Medicine. https://www.fda.gov/medical-devices/in-vitro-diagnostics/precision-medicine (accessed 19 April 2024).

United States Food and Drug Administration (2018b). Step 3: Clinical Research. https://www.fda.gov/patients/drug-development-process/step-3-clinical-research (accessed 28 January 2024).

United States Food and Drug Administration (2019). Adaptive Designs for Clinical Trials of Drugs and Biologics: Guidance for Industry. https://www.fda.gov/media/78495/download (accessed 01 Dec 2022).

United States Food and Drug Administration (2021). Adjusting for Covariates in Randomized Clinical Trials for Drugs and Biologics with Continuous Outcomes: Guidance for Industry. https://www.fda.gov/media/148910/download (accessed 12 May 2022).

United States Food and Drug Administration (2023). Optimizing the Dosage of Human Prescription Drugs and Biological Products for the Treatment of Oncologic Diseases. https://www.fda.gov/media/164555/download (accessed 11 April 2024).

Van Lancker, K., Betz, J., and Rosenblum, M. (2022). Combining Covariate Adjustment with Group Sequential and Information Adaptive Designs to Improve Randomized Trial Efficiency. *arXiv preprint arXiv:2201.12921*.

Villar, S. and Rosenberger, W. (2018). Covariate-Adjusted Response-Adaptive Randomization for Multi-Arm Clinical Trials Using a Modified Forward Looking Gittins Index Rule. *Biometrics*, 74(1):49–57.

Wages, N. A. and Conaway, M. R. (2014). Phase I/II Adaptive Design for Drug Combination Oncology Trials. *Statistics in Medicine*, 33(12):1990–2003.

Walter, S., Guyatt, G., Bassler, D., Briel, M., Ramsay, T., and Han, H. (2019). Randomised Trials with Provision for Early Stopping for Benefit (or Harm): the Impact on the Estimated Treatment Effect. *Statistics in Medicine*, 38(14):2524–2543.

Wang, K. and Ivanova, A. (2005). Two-Dimensional Dose Finding in Discrete Dose Space. *Biometrics*, 61(1):217–222.

Wang, W. and Yan, J. (2021). Shape-Restricted Regression Splines with R Package splines2. *Journal of Data Science*, 19(3):498–517.

Wang, Y. and Ivanova, A. (2015). Dose Finding with Continuous Outcome in Phase I Oncology Trials. *Pharmaceutical Statistics*, 14(2):102–107.

Wang, Z. and de Freitas, N. (2014). Theoretical Analysis of Bayesian Optimisation with Unknown Gaussian Process Hyper-Parameters. *arXiv preprint arXiv:1406.7758*.

Wang, Z., Zhang, J., Xia, T., He, R., and Yan, F. (2023). A Bayesian Phase I-II Clinical Trial Design to Find the Biological Optimal Dose on Drug Combination. *Journal of Biopharmaceutical Statistics*. Epub ahead of print.

Wason, J. M., Mander, A. P., and Eisen, T. G. (2011). Reducing Sample Sizes in Two-Stage

Phase II Cancer Trials by Using Continuous Tumour Shrinkage End-Points. *European Journal of Cancer*, 47(7):983–989.

Willard, J., Golchi, S., and Moodie, E. E. (2024). Covariate Adjustment in Bayesian Adaptive Randomized Controlled Trials. *Statistical Methods in Medical Research*, 33(3):480–497.

Willard, J., Golchi, S., Moodie, E. E., Boulanger, B., and Carlin, B. P. (2023). Bayesian Optimization for Personalized Dose-Finding Trials with Combination Therapies. *arXiv preprint arXiv:2310.17334*.

Williams, C. K. and Rasmussen, C. E. (2006). *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, Massachusetts.

Woodcock, J. and LaVange, L. M. (2017). Master Protocols to Study Multiple Therapies, Multiple Diseases, or Both. *New England Journal of Medicine*, 377(1):62–70.

Yuan, Y., Hess, K. R., Hilsenbeck, S. G., and Gilbert, M. R. (2016a). Bayesian Optimal Interval Design: A Simple and Well-Performing Design for Phase I Oncology Trials. *Clinical Cancer Research*, 22(17):4291–4301.

Yuan, Y., Lee, J. J., and Hilsenbeck, S. G. (2019). Model-Assisted Designs for Early-Phase Clinical Trials: Simplicity Meets Superiority. *Journal of Clinical Oncology, Precision Oncology*, 3:1–12.

Yuan, Y., Nguyen, H. Q., and Thall, P. F. (2016b). *Bayesian Designs for Phase I-II Clinical Trials*. CRC Press, New York, New York.

Yuan, Z., Chappell, R., and Bailey, H. (2007). The Continual Reassessment Method for Multiple Toxicity Grades: A Bayesian Quasi-Likelihood Approach. *Biometrics*, 63(1):173–179.

Zhang, J., Lin, R., Chen, X., and Yan, F. (2024). Adaptive Bayesian Information Borrowing

Methods for Finding and Optimizing Subgroup-Specific Doses. *Clinical Trials*. Epub ahead of print.

Zhilinskas, A. (1975). Single-Step Bayesian Search Method for an Extremum of Functions of a Single Variable. *Cybernetics and Systems Analysis*, 11(1):160–166.