CONSIDERING DOSE REDUCTIONS AND DOSE INTERRUPTIONS IN ONCOLOGY DOSE-FINDING TRIAL DESIGN

Jesse Elijah Helman

Chambersburg, Pennsylvania

Bachelor of Science, Shippensburg University, 2017

A Dissertation submitted to the Graduate Faculty

of the University of Virginia in Candidacy for the Degree of

Doctor of Philosophy

Department of Statistics

May 2023

Dr. Karen Kafadar, Chair

Dr. Mark Conaway

Dr. Xiwei Tang

Dr. Steve Boker

Considering dose reductions and dose interruptions in oncology dose-finding trial design

Jesse Elijah Helman

(ABSTRACT)

Within the past decade, the FDA has issued many post-marketing requirements or commitments to study alternate doses in Molecularly Targeted Agents and Immunotherapies in reaction to a large percentage of Dose Reductions and Dose Interruptions, caused by chronic low-grade toxicities. Almost all of these trials used the Maximum Tolerated Dose (MTD) based on Dose Limiting Toxicity (DLT) information only. While extensions of traditional cytotoxic dose-finding designs address long-term toxicities as well as chronic low-grade toxicities, none has explicitly and appropriately estimated the probability of Dose Reduction and Dose Interruption. This dissertation introduces a novel modeling procedure to model Dose Reductions (and/or Dose Interruptions) and DLTs simultaneously, adopting a semi-competing risks framework for the first time in dose-finding. Simulations are presented to evaluate operating characteristics and design options, and to compare the proposed method to others likely to be used in the same setting. Important quantities and some theoretical results are derived. Additionally, Dose Reductions and DLTs affect the estimation of efficacy, which is necessary to evaluate in the current setting. No previous methods have addressed Dose Reductions in the context of early phase efficacy. A new efficacy estimand is defined for a tumor burden outcome to appropriately handle Dose Reductions and DLTs, based on Multiple Imputation. An illustration and discussion compare the proposed estimand with current practice. Finally, dose selection options are discussed using toxicity and efficacy information.

Acknowledgments

I would like to express my gratitude to my advisor Dr. Mark Conaway. His motivation and guidance made this process truly enjoyable. Similarly, to my co-advisor Dr. Karen Kafadar, her contributions to this dissertation were invaluable. I would like to thank my committee members, Dr. Xiwei Tang and Dr. Steve Boker, for their service and contribution.

I would like to acknowledge my summer internship advisors, specifically Karin Bowen, Ales Kotalik, and Patrick Darken at AstraZeneca who introduced me to clinical estimands, the subject of a large piece of this work. I would also like to thank my parents, Carl and Rebekah Helman, for their support, and my friend and colleague Noah Gade for countless hours of discussion and brainstorming that provided the foundation for the ideas presented in this dissertation.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Traditional Phase I trials in oncology consider toxicity.

Traditionally, the major objective of a dose-finding clinical trial in oncology for cytotoxic therapy (chemotherapy) is to identify the highest dose that can be given safely in humans (Maximum Tolerated Dose), which will be evaluated for efficacy in subsequent Phase II and Phase III trials. The Maximum Tolerated Dose (MTD) is defined as the dose x^* at which the probability of a Dose Limiting Toxicity (DLT) Y equals a pre-specified limit θ,

$$P(Y = 1|x^*) = \theta \tag{1.1}$$

where the DLT endpoint Y is binary. A DLT is one of a collection of possible toxicities defined before a trial as intolerable, typically a Grade 3 or higher on the National Cancer Institute *Common Terminology Criteria for Adverse Events (CTCAE)* 2023 scale.

Contemporary dose-finding trials may also consider efficacy with toxicity.

The increasing prevalence of Molecularly Targeted Agents (MTA) and Immunotherapies means the majority of dose-finding trials are investigating therapies where the toxicity and efficacy mechanisms are separate. Under the traditional paradigm of cytotoxics, the same mechanism generates toxicity and efficacy. It is then only necessary in dose-finding to set a limit on the acceptable toxicity level (MTD). Conversely,

in the paradigm of the MTA and immunotherapy, the mechanism of action is not cytotoxic. Efficacy may not be a function of toxicity. Toxicities are truly side effects, and higher doses may not lead to greater efficacy. Therefore, it is not enough to identify the MTD. Toxicity and efficacy must each be modeled, and the goal of dose-finding is to find the Optimal Biological Dose (OBD), defined as the dose with acceptable toxicity that maximizes efficacious response. Moreover, drug developers have an added motivation to accelerate drug trials and reduce costs. Evaluation of efficacy from the start can help integrate previously separate trials, reducing time and money spent in development.

MTAs and immunotherapies also differ from the traditional oncology paradigm in their toxicity profile. Immunotherapies are given chronically, typically over a 6-24 month window or longer, and toxicities are observed beyond the one-month window used to determine the MTD in cytotoxic therapies. Postel-Vinay et al. 2014 analyzed 54 completed Phase I trials of MTAs given as monotherapies and found 599 of 2,084 patients had a Grade 3 toxicity or higher, 50% of which occurred after the first month. Similarly, 53% of 1,242 patients with a Grade 2 toxicity or higher happened after the first month.

Chronic Grade 2 toxicities turn into Dose Reductions.

The protocol in oncology trials aims to keep patients on the trial at their originally assigned dose until they complete the entire course of treatment or experience a DLT, at which point they are taken off the therapy and the trial. However, the Grade 2 toxicities experienced on MTAs and immunotherapies are often chronic, and can significantly affect a patient's quality of life, yet do not fit into the definition of a DLT. Thus, protocols have evolved to "manage" these chronic toxicities by reducing the patient's dose from its initial level, allowing the chronic toxicity to resolve,

while also keeping the patient on the trial (Postel-Vinay et al. 2014). This is a Dose Reduction. However, in analyzing these trials, Dose Reductions are not incorporated when estimating toxicity and efficacy over a dose range.

Dose Reductions are often the result of patient-reported outcomes, a topic of expanding interest in all clinical trials. Patient-reported outcomes allow a patient to initiate the evaluation of toxicities and change the frame of reference for what constitutes an intolerable toxicity, rather than be driven solely by a clinician applying toxicity criteria uniformly to all patients. A patient may determine a chronic low-grade toxicity is intolerable, one that had not previously risen to the level of DLT evaluated at clinician visits. By reporting such a toxicity, the patient would then have their dose interrupted or reduced by a clinician in an attempt to resolve the toxicity. Thus, the scope of this dissertation can project into a much larger discussion of patient-reported outcomes in clinical trials.

Ignoring Dose Reductions has substantial clinical implications.

Ignoring chronic Grade 2 toxicities and resulting Dose Reductions must underestimate relevant toxicity. Thus, the chosen Phase II or Phase III dose may be too high, too toxic. This is evidenced in practice. Through 2016, 31 small-molecule kinase inhibitors (KI), a specific class of MTAs, had been approved by the FDA for oncology indications. However, among these 31 approvals, eight contained post marketing requirements (PMR) or commitments (PMC) to study alternate doses, as the FDA believed that the optimal dose may not have been identified because of the large percentage of Dose Reductions among patients (Janne et al. 2016). Each of the therapies in Figure 1.1, save for Idelalisib, used the MTD as the dose for its initial registration trial (Bullock, Rahman, and Liu 2016).

Table 1. Dose interruptions and reductions in initial registration trials for small-molecule KIs approved for oncology indications with PMC or PMR to study alternate doses (percentage of patients on registration studies)

Drug	Dose interruption	Dose reduction	Dose interruption or delay
Erlotinib	62%	19%	NA
Vandetanib	47%	49%	80%
Cabozantinib	NA	79%	86%
Ponatinib	66%	52%	74%
Ceritinib	69%	59%	71%
Idelalisib	NA	34%	53%
Lenvatinib	56%	68%	90%

Table 1.1: Dose Interruptions and Dose Reductions for approved small molecule KIs with PMC or PMR to study alternate doses from Janne et al. 2016.

The FDA is giving a clear signal that Dose Reductions relay relevant information about toxicity and will affect the outcome of a drug's application. Additionally, the FDA–AACR Oncology Dose-Finding Workshop (Part 3) in 2017, which focused on immunotherapies and best practices regarding patient and dose selection, reported that "Dose Reduction from the maximum-tolerated dose (MTD) or recommended Phase II dose for molecularly targeted agents should be considered for the starting dose, due to the possibility of unexpected toxicities" (Emens et al. 2017). Note that the concept of "unexpected" is applicable only in the typical analysis framework. If Dose Reductions were considered along with DLTs during an appropriate time window, these "unexpected" toxicities would instead be observed or estimated. Consequently, one research question becomes clear: How do we properly model dose-toxicity to properly include and account for Dose Reductions?

Proposed method accounts for Dose Reductions in estimating toxicity.

In this dissertation, a method is introduced to model Dose Reduction toxicities (DRs) and DLTs simultaneously, adopting a semi-competing risks framework for the first time in dose-finding. A DLT can be seen as a terminal event, causing patient removal from a trial; whereas, a DR is a non-terminal event, allowing continuation on a trial. Observation of a DR is subject to the observation of a DLT, though, thus defining a semi-competing risks framework. Conditional models for the hazard functions of the non-terminal and terminal events are constructed in the spirit of Lee et al. 2015 and Putter, Fiocco, and Geskus 2007. As patients are entered sequentially in early-stage dose-finding trials, and the ethical practice is to use the information gathered from previous patients to inform how to best dose new patients, a Bayesian framework is utilized to continuously update estimates. A Bayesian estimation procedure for model parameters is detailed, and closed-form solutions are presented for calculating relevant probabilities regarding the two events.

Dose Reductions also affect the estimation of efficacy.

As previously stated, proper dose-finding for immunotherapies in oncology should model toxicity and efficacy, and Dose Reductions can have a significant impact on efficacy as well. According to Sachs et al. 2016, overshooting the MTD by ignoring Dose Reductions "can have significant impact on efficacy, as it has resulted in the use of less tolerable doses over longer duration when lower doses would suffice, thus causing poorer adherence and concomitant lower efficacy." Beyond adherence, Dose Reductions confuse how to model dose-efficacy. If a patient is enrolled in a trial at an initial dose and then relatively soon after dose reduced, no obvious procedure is established regarding how to associate this patient's efficacy profile with the doses received. From Janne et al. 2016, the "big problem is unknown efficacy at each dose

level because patients are dose reduced so often and efficacy data is not reliable if reported as the starting dose. ...unknown efficacy in light of frequent Dose Reductions in the postmarket setting raises the question of whether efficacy reported in early-phase trials is accurate when applied to a real-world population." Thus, a second research question becomes clear, similar to the first: How can dose-efficacy be appropriately modeled incorporating Dose Reductions to make dosing decisions?

Dose Reductions in a continuous longitudinal efficacy response.

This dissertation discusses modeling efficacy based upon a continuous longitudinal outcome. An example common in Phase I and Phase II trials is percent change-from-baseline in tumor burden. DLTs and DRs complicate the estimation of this outcome. They are intercurrent events, using the language of estimands, events which either preclude the observation of the outcome or affect its measurement or interpretation. This becomes a problem of identifying the proper estimand. Ideally, patients are observed for the full period on their initial dose, offering complete change-from-baseline data. However, DLTs cause removal from the trial which induces missing data. The question arises, should only observed data be used, effectively changing the time scale of the endpoint by patient, or should data be imputed for the original endpoint?

The complexities of DRs are more subtle. A Dose Reduction midway through follow-up will almost surely reduce the efficacy of the treatment, compared to remaining on the initial dose until complete follow-up. But, another question arises, does this matter? Should the negative contribution of a Dose Reduction be "separated" from the hypothetical complete outcome on initial dose? With the example of percent change-from-baseline in tumor burden, issues of monotonicity also play a role in deciding what estimands are appropriate. Ultimately, an estimand is proposed

that imputes missing data based on what is expected to happen to a patient after a DLT, including the occurrence of Dose Reduction in the imputation model. Finally, options for dose selection are presented that combine estimates from the toxicity and efficacy models.

1.1 Literature Review

First, common methods to estimate the MTD in a Phase I trial are reviewed based strictly on toxicity data as well as how efficacy has been retroactively built into these methods. Dosing decisions in realization of the objectives of a Phase I trial historically have been made using rule-based designs, such as the 3+3 design, which are still common today. However, rule-based designs are now appropriately being pushed aside in favor of modeling methods. Rule-based designs will not be discussed here.

Traditional Toxicity Methods

One of the first statistical model-based designs for dose allocation and dose decision was the Continual Reassessment Method (CRM) by O'Quigley and Shen 1996. As with other similarly-aimed model-based designs, CRM uses a Bayesian framework to sequentially update and estimate the dose level at which to treat the next available patient based on binary toxicity data (DLT). The CRM is developed under the assumption of a discrete set of doses selected before a study.

First, some functional dose-toxicity curve, denoted by $F(x_i, \beta)$, is considered to model the probability of a DLT. The CRM assumes that this function is monotonic increasing in dose x_i and fixed parameter β. As the dose increases, the probability of a DLT is non-decreasing for a fixed parameter value; as the parameter value increases,

the probability of a DLT is non-decreasing for a fixed dose level. The function $F(x_i, \beta)$ models the true DLT probabilities $\pi(x_i) = P(Y_j = 1|x_i)$ at each given dose i. One example of a function F, most relevant to this paper, is the basic power model, where $F(x_i, \beta) = x_i^{exp(\beta)}$. The basic power model is a single parameter model, able to build in the assumption of monotonic increasing probability with dose by applying it to the "skeleton" without having to use an extra parameter.

The Bayesian framework places a prior distribution $h(\beta)$ on the parameter β. Once $\hat{\beta}_j$ is estimated using the posterior distribution, it can be used to find an estimated probability of DLT at each dose x_i in the discrete dose set, $\hat{\pi}(x_i) = \hat{F}(x_i, \beta_j)$. Some measure of distance, such as a simple Euclidean (L^1) distance, can be used to find the dose that has an estimated probability of toxicity closest to the target; e.g. $x_i = argmin_i|\hat{F}(x_i, \beta_j) - \theta|$.

The basic CRM, as explained previously, addresses the problem of choosing a single MTD from a fixed, discrete set of doses of one cytotoxic drug. Further methods that make MTD decisions from a fixed, discrete dose set include Thall et al. 2003, Wang and Ivanova 2005, Yin and Yuan 2009, and Braun and Wang 2010, among many others. Each method builds on the CRM in various ways.

Efficacy Additions into Traditional Toxicity Methods

Associated literature contains extensions to the CRM where efficacy information is considered in addition to toxicity information when selecting a dose for further study, to accommodate the MTA and immunotherapy paradigm. Mostly, efficacy has been treated as a binary variable similar to toxicity. The probability of efficacy at dose level x_i for patient j can be represented by $\pi_E(x_i) = P(Z_j = 1|x_i)$ where Z is a binary indicator of efficacy. Typically, Z would be dichotomized from a continuous measure of efficacy such as tumor shrinkage or CD8 cell counts. The statistic from

Z commonly referenced and analyzed is the "response rate," defined as the rate of patients on a given dose who exhibit an efficacious or antidisease response above a threshold, say θ. Instead of exclusively aiming to limit probability of DLT to a target level, $P(Y_j = 1|x_i) = \theta$, the goal is to remain within a safe probability of toxicity while maximizing the probability of efficacy $P(Z_j = 1|x_i)$. In this way, the dose of choice is better referred to as the Optimal Biological Dose (OBD) rather than the MTD. Wages and Tait 2015 developed a method to address this problem built on the foundation of the CRM under the fixed, discrete dose set paradigm. Their method assumes

$$\pi_E(x_1) \leq ... \leq \pi_E(x_\nu) \geq ... \geq \pi_E(x_K), \tag{1.2}$$

where the dose x_ν is the dose closest to the maximum probability of efficacy under a unimodal dose-efficacy relationship or the beginning of the plateau in such a dose-efficacy relationship. Possible orderings of dose-efficacy fitting (1.2) are enumerated and treated as separate models $q = 1, ..., M$, and model selection is used to choose the best ordering based on the data. The model is subsequently fit via CRM based on the chosen ordering.

The literature discussed thus far provided the entry point into methods considering both toxicity and efficacy, but the context of interest clearly differs from that discussed above. For the current context, toxicity involves two separate binary events instead of one, both evaluated beyond Cycle 1, while efficacy is considered ultimately in its original form, as a longitudinal continuous measure, the full measure of efficacy, versus a single point in time or a maximum gained effect. Literature relating to one or more of these differences is discussed in the following.

Long-term Toxicities

Long-term toxicities alone are addressed multiple ways in the literature. Under a long-term observation window, a patient may be in the middle of their follow-up period when a dosing decision needs to be made for a new patient coming onto the trial. Thus, the former patient can be seen as censored. The Time-to-Event CRM method or TITECRM (Cheung and Chappell 2000) weights the likelihood of subjects who have not experienced toxicity (censored observations) by the length of their follow-up times at the time of analysis. Similar weighting mechanisms are utilized in Braun 2006, Mauguen, Deley, and Zohar 2011, and Lin et al. 2016. Imputing censored observations in a Bayesian framework, known as Bayesian Data Augmentation, has been used to address long-term toxicities by Liu and Ning 2013 and Liu, Yin, and Yuan 2013. Yuan and Yin 2011 use an Expectation Maximization (EM) method to estimate dose-toxicity with censored observations. Yin and Yuan 2009 model toxicity and efficacy as time-to-event outcomes in a traditional Cox multiplicative hazards survival framework, selecting dose based on the ratio of the area under survival curve for toxicity and efficacy. However, none of these methods addresses multiple toxic outcomes or Dose Reductions.

Multiple Toxicities

The CRM Multiple Constraints (CRM-MC) method by Lee, Cheng, and Cheung 2011 allows for the specification of multiple toxicity thresholds to address multiple possible toxic outcomes, including "moderate toxicities," the chronic low-grade toxicities that are responsible for Dose Reductions. The CRM is used for each toxicity, and the chosen dose is the largest dose with estimated toxicity probabilities under each threshold. Some methods address multiple possible toxic outcomes by treat-

ing toxicity as an ordinal variable. Paoletti et al. 2015 and Meter, Garrett-Mayer, and Bandyopadhyay 2012 use cumulative logit formulations to estimate probability of each toxic outcome, the latter in an explicit CRM framework. Ivanova and Kim 2019 present a broad approach that can be used for ordinal outcomes assuming an underlying continuous objective function. Lee et al. 2012, among many others, develop a Toxicity Burden Score to measure total toxicity over time. However, these methods are not developed for the long-term toxicity problem or for the dose change that happens when a DR occurs.

One simple method that considers Dose Reductions is the longitudinal Relative Dose Intensity (RDI) procedure proposed by Hirakawa, Yonemori, et al. 2018 and Hirakawa, Tanaka, and Kaneko 2019. In this design, the average dose a patient receives over the entire follow-up period is their mRDI (mean relative dose intensity). If a patient experiences a DLT, the dose is assumed to be zero for the rest of the prescribed follow-up period. Thus, if a patient experiences a DLT early, the mRDI will be very low, whether or not Dose Reductions were experienced before the DLT. The mRDI is averaged over all patients who started at the same dose; this value is the pRDI (population Relative Dose Intensity) for that dose. A dose is considered acceptable if the pRDI is greater than 0.75.

TITECRM Multiple Constraints

The only existing design in the literature that considers the same multiple toxicity, long-term context is the TITECRM-MC method from Lee et al. 2019. The authors combine TITECRM and CRM-MC to estimate the probability of multiple toxicities over a dose range allowing censored data. Let Z be an ordinal toxicity outcome, where $Z = 0$ if the patient does not experience a toxicity; $Z = 1$ if the patient experiences moderate toxicity without a DLT; $Z = 2$ if the patient experiences a

DLT. The target probabilities are defined as π_{MT} and π_{DLT}; thus, the chosen dose θ is the maximum dose such that $P(Z \geq 1|x) \leq \pi_{MT}$ and $P(Z \geq 2|x) \leq \pi_{DLT}$,

$$\theta = argmax_x P(Z \geq t_l|x) \leq \pi_l, l = 1, 2. \tag{1.3}$$

As in the original CRM, a working model is assumed $P(Z \geq t_l\ x) = F_l(x; \beta)$, this time with two parameters $\beta = (\beta_1, \beta_2)^T$ relating to the two types of toxicity.

Semi-Competing Risks

The semi-competing risks framework used in this paper for modeling toxicity was not developed, nor has been proposed, for dose-finding scenarios. Lee et al. 2015 worked in the context of pancreatic cancer, with death as the terminal outcome, and readmission to the hospital following initial diagnosis as the non-terminal outcome. In semi-competing risks, the main challenge is the unidentifiability of the marginal survivor function of the non-terminal outcome because occurrence of the terminal outcome excludes occurrence of the non-terminal outcome. The marginal survivor function for the terminal outcome can be fully observed. From Lee et al. 2015, methods to deal with this unidentifiability typically fall into two groups: modeling the dependence with a copula or building conditional hazard functions for the terminal and non-terminal events.

Lee et al. 2015 take the path of conditional hazard functions. Let T_1 represent the time to non-terminal event (Dose Reduction), and T_2 represent the time to terminal event (DLT), both taken from entry into the trial. A cause-specific hazard is built for the non-terminal event, $h_1(t_1)$, for the terminal event, $h_2(t_2)$, and for the terminal event conditional on the previous occurrence of the non-terminal event, $h_3(t_2|t_1)$. The hazard functions are defined on $0 < t_1 < t_2$ in a Cox multiplicative hazards style as

follows,

$$h_1(t_{1i}|\gamma_i, x_i) = \gamma_i h_{01}(t_{1i}) e^{x_i^T \beta_1}, t_{1i} > 0; \tag{1.4}$$

$$h_2(t_{2i}|\gamma_i, x_i) = \gamma_i h_{02}(t_{2i}) e^{x_i^T \beta_2}, t_{2i} > 0; \tag{1.5}$$

$$h_3(t_{2i}|t_{1i}, \gamma_i, x_i) = \gamma_i h_{03}(t_{2i}) e^{x_i^T \beta_3}, 0 < t_{1i} < t_{2i}, \tag{1.6}$$

where h_{0g} $(g = 1, 2, 3)$ is a baseline hazard function; x_i is a $p \times 1$ vector of covariates; β_g is a vector of p log-hazard ratio regression parameters; γ_i is a shared frailty (random effect, taken to be distributed independent of x_i). If a patient experiences the terminal event prior to the non-terminal event, then T_1 is set to ∞, and the remaining probability mass not contained in $0 < T_1 < T_2$ is sent to the line $T_1 = \infty$.

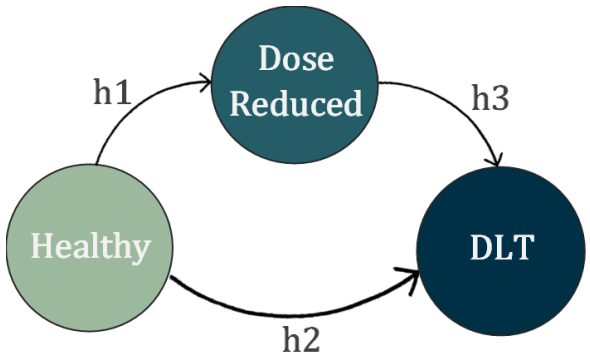

Figure 1.1: Illness-death (three-transition) semi-competing risks model with Dose Reductions and DLTs.

The method developed in this paper uses the conditional hazard function formulation. An Accelerated Failure Time formulation can be seen as a close relative (Lee, Rondeau, and Haneuse 2017). This formulation of the semi-competing risks framework can also be viewed as a multi-state model, as in Putter, Fiocco, and Geskus 2007 and Hougaard 1999. The states "Healthy," "Dose Reduced," and "DLT," define

an illness-death model with interest lying in the probability of transition between the states at a given time. A diagram of the illness-death multi-state model in the current setting is shown in Figure 1.1.

Estimands

Discussion of efficacy modeling in this dissertation uses the language of estimands. The Estimand framework, resulting from the International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use (ICH) E9 guideline addendum, expands on the definition of a statistical estimand (target parameter to be estimated) to as well list possible intercurrent events and how they will be handled (ICH 2017). It is now expected to be detailed in a separate section of a study protocol.

Estimand Strategy	Definition
Treatment Policy	Intention-to-treat. Measure value of the outcome of interest as observed, regardless of the occurrence of the intercurrent event.
Composite	Intercurrent event included in the definition of outcome variable.
Hypothetical	Assumes and attempts to estimate as if the intercurrent event did not happen.
Principal Stratum	Measurement of variable of interest limited to a subgroup unlikely to experience intercurrent events.
While-on-Treatment	Measurement until the time of event and not after.

Table 1.2: Estimand strategies to handle intercurrent events.

Intercurrent events, as stated previously, are events that either preclude the observation of the outcome or affect its measurement or interpretation. Some examples of intercurrent events are: discontinuation of treatment due to lack of efficacy, discontinuation of treatment due to toxicity, surgery related/unrelated to the study,

starting alternative treatment while on study, developing a side-effect, and experiencing a treatment change but continuing treatment. Each of these events affects the measurement of the true effect of therapy. Five main strategies address intercurrent events, explained in Table 1.2. If multiple intercurrent events are considered, different estimand strategies could be used for each in combination.

1.2 Motivating Trials

This work is motivated by considerations of real trials that involve MTAs and immunotherapies as well as the practical biologic foundations of these therapies. One such trial is LURET, an open-label, multi-center, Phase II trial of vandetanib, a multitargeted tyrosine kinase inhibitor exhibiting RET kinase activity (a class of MTA), for patients with advanced RET-rearranged non-small cell lung cancer (Yoh et al. 2017). All patients in the trial were initially assigned to receive 300 mg of oral vandetanib daily. Therapy was continued until disease progression, unacceptable toxicity, death, or withdrawal from the study. Toxicities were graded according to the CTCAE version 4.03. Tumor burden was evaluated based on Response Evaluation Criteria in Solid Tumors (RECIST) version 1.1 (Eisenhauer et al. 2009). Tumor response was assessed every four weeks. Although this trial was Phase II and considered a single dose, it is illustrative of the endpoints considered for both toxicity and efficacy that would ideally be used in a Phase I/II design considering multiple doses.

Of the 19 patients eligible for safety analysis in LURET, four had a DLT and were removed from the trial; 16 patients had a Dose Interruption due to an adverse event; and 10 had a Dose Reduction (to 200 mg) due to an adverse event. Thus, with only four DLTs, daily 300 mg vandetanib was deemed to be "tolerable" with a "manageable safety profile," as is common. The timing of Dose Reductions for each patient

in the study was not published. The primary efficacy analysis evaluated whether a patient had an "objective response" or not. Objective response was a dichotomization of at least 30% maximum tumor shrinkage achieved during the follow-up period taken from the full longitudinal tumor shrinkage profiles presented in Figure 1.2 (sustained for at least four weeks). The RECIST guidelines classified nine of the 17 patients in Figure 1.2 as objective responders.

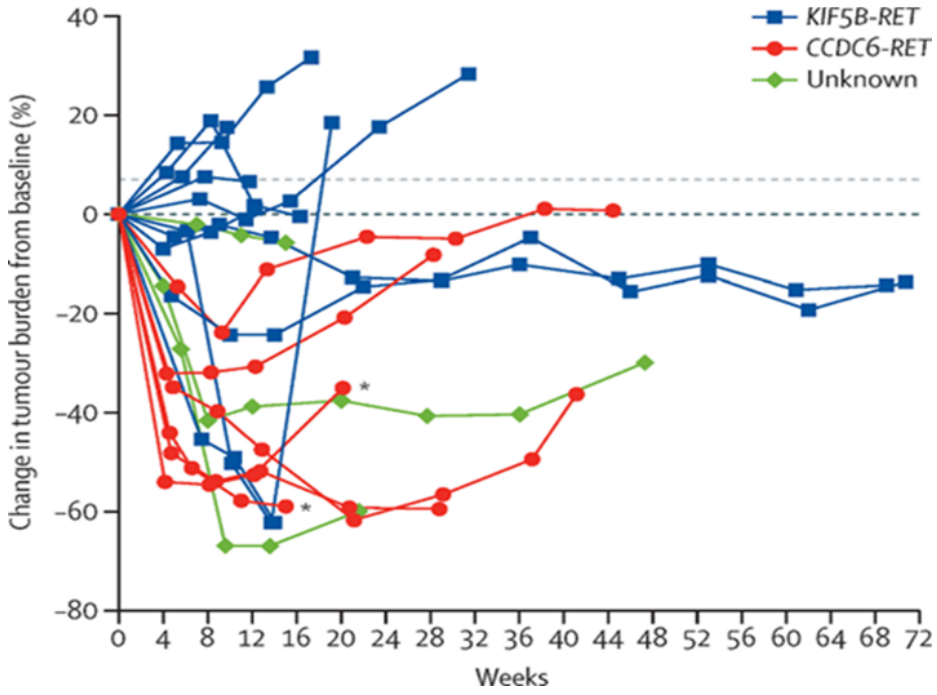

Figure 1.2: Percent change in tumor burden from baseline, from Yoh et al. 2017.

Table 1.1 presents other examples of therapies that caused a significant percentage of Dose Reductions. More such examples can be found in Roda, Jimenez, and Banerji 2016. In regards to efficacy, rarely does a trial publish longitudinal tumor response data. Rather, these trials take some maximum change-from-baseline tumor shrinkage achieved for each patient and use a responder analysis to categorize efficacy outcomes. The few other examples of published longitudinal tumor shrinkage data

that could be found are presented in Appendix A.1: Wolchok et al. 2013, Seiwert et al. 2016, and Hamid et al. 2013.

Chapter 2

Toxicity Methods

A basic illness-death model contains three possible transitions. In the context of Dose Reductions and DLTs, these transitions are: (1) Healthy to Dose Reduced; (2) Healthy to DLT; (3) Dose Reduced to DLT. A Healthy state simply represents the state of the patient at randomization. It is possible, though rare, in immunotherapy trials to be Dose Reduced twice (Johnson et al. 2005, Philip et al. 2005). In this case, there are five possible transitions: the three previous transitions; (4) Dose Reduced once to Dose Reduced twice; (5) Dose Reduced twice to DLT. These states and transitions are diagrammed in Figure 2.1.

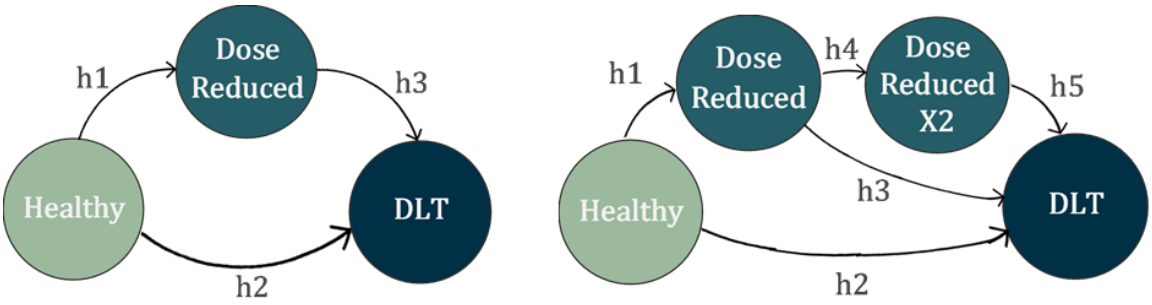

Figure 2.1: Illness-death and five-transition semi-competing risks models with Dose Reductions and DLTs.

In the illness-death model, a DLT can be reached via two unique paths. In the five-transition model, there are three unique DLT paths instead of two, and the "Dose Reduced Twice" state cannot be reached directly from the "Healthy" state.

Depending on the protocol of the trial, the illness-death or five-transition model may be more appropriate. Far more common are trials that allow and observe only a single possible Dose Reduction for each patient. Details of the illness-death model are specified here as well as in Chapter 3, with details of the five-transition model referred to Appendix B.

2.1 The Three-Transition Toxicity Model

Let T_1 be the time to DR and T_2 be the time to DLT, measured from entry on the trial. The hazards in limit form can be written as follows:

$$h_1(t_1) = \lim_{\Delta \to 0} \frac{P(T_1 \in [t_1, t_1 + \Delta) | T_1 \geq t_1, T_2 \geq t_1)}{\Delta} \tag{2.1}$$

$$h_2(t_2) = \lim_{\Delta \to 0} \frac{P(T_2 \in [t_2, t_2 + \Delta) | T_1 \geq t_2, T_2 \geq t_2)}{\Delta} \tag{2.2}$$

$$h_3(t_2|t_1) = \lim_{\Delta \to 0} \frac{P(T_2 \in [t_2, t_2 + \Delta) | T_1 = t_1, T_2 \geq t_2)}{\Delta} \tag{2.3}$$

where $h_1(t_1)$ represents transition 1→2; $h_2(t_2)$ represents the transition 1→3; and $h_3(t_2|t_1)$ represents the transition 2→3. The hazard functions in Lee et al. 2015 are given a Cox-style formulation with a shared frailty, $h(x|\mathbf{z}) = \gamma_i h_0(x) exp\ (\beta^t \mathbf{z})$, where $h_0(x)$ is a baseline hazard rate and z is a vector of covariates. The shared frailty is analogous to the random intercept used in mixed models to induce dependence among clustered responses.

Constant-Skeleton Specification

However, the assumption in dose-finding is a small data setting. Thus, the power model from the CRM, with dose on the "skeleton," is adapted in specifying the hazard functions. The hazard functions remain multiplicative, but do not have the Cox-style

baseline hazards model formulation. The shared frailty γ_i is foregone for the same small data reasoning. This formulation assumes no dependence between T_1 and T_2, which rigorously is incorrect, but valuable in the desire to estimate one less parameter that would not anyway provide much added benefit for the prediction tasks of interest. The simplest formulation assumes constant baseline hazards λ_1 and λ_2 for T_1 and T_2, respectively. Let i represent one of a sample of n patients:

$$h_1(t_1|x_j) = \lambda_1 x_{ij} \tag{2.4}$$

$$h_2(t_2|x_j) = \lambda_2 x_{ij} \tag{2.5}$$

$$h_3(t_2|t_1, x_j) = \lambda_2 x_{i,j-1}^{exp(\beta_1)} \tag{2.6}$$

This is a three-parameter model, $\lambda_1, \lambda_2, \beta_1$. The parameter $exp\ (\beta_1)$ is a "history" parameter. No assumptions are made about the sign of β_1; if β_1 is positive, the hazard of a DLT is assumed to increase after a patient experiences a DR, given the dose to which they have been reduced. This is analogous to saying a DR puts a patient in a "weakened" position, more susceptible to toxicity. If β_1 is negative, the hazard of a DLT is assumed to decrease after a patient experiences a DR, given the dose to which they have been reduced. This is analogous to saying a DR has induced resistance to the therapy, causing the patient to become less susceptible to toxicity. If clinicians have information regarding this assumption before the trial, it can be incorporated in the prior for β_1.

The doses x_j are scaled to the $[0, 1]$ interval, remaining in their original ordering assuming monotonic increasing toxicity with dose. Notice that a reduced dose, x_{j-1}, is included in h_3. This is a fundamental difference from TITECRM-MC, where the initial dose is used exclusively in all modeling. The proposed model assumes the reduced dose is the cause of a DLT that occurs after a DR, not the initial dose, and

should more appropriately represent the toxicity of each dose. This assumption also guides the assessment of efficacy in Chapter 5.

Baseline hazards cannot be left unspecified in a Bayesian setting. The baseline hazards are initially assumed to be constant, which may be an oversimplification, but is motivated by the sparse-data setting. The baseline hazard rate of a DR is λ_1, and the baseline hazard rate of a DLT is λ_2, assuming $0 < \lambda_2 \leq \lambda_1$.

Hypothetically, the parameter $exp\ (\beta_1)$ could be omitted as well, creating a two-parameter model. Then, the hazard of a DLT before a DR (h_2) compared to the hazard of a DLT after a DR (h_3) would be entirely determined by the scaling of x_j and x_{j-1}. This places large importance on the choice of skeleton; consequently, $exp\ (\beta_1)$ allows some valuable flexibility. A Cox-style with a frailty term would involve five parameters to estimate. Modeling in dose-finding designs is always a balance between necessary complexity and valuable parsimony due to limited data available for estimating large, more complex models, with the balance often favoring valuable parsimony.

Notice that h_3, the conditional hazard of DLT given DR, does not depend on the time of DR, t_{1i}. This is labeled in the literature as a "Markov" model. Conversely, if h_3 were modeled using the time since DR $(t_{2i} - t_{1i})$, the model would be labeled as "Semi-Markov." The five-transition model hazard rates can be found in Appendix B.

Weibull-Skeleton Specification

Though long-term toxicities are expected with targeted agents and immunotherapies, it may not be reasonable to assume constant hazards over the entire follow-up period. Rather, it may be that most toxicities are observed early, within a few weeks, or that most are a result of cumulative toxicities and occur later in the observation window.

If observed data truly follow a non-constant hazard form, then fitting a constant hazard model form can greatly bias the dose-selection process, given the amount of partial follow-up expected. Specifically, if hazard rates are truly decreasing and a sizable percentage of patients have not completed follow-up (partial information), then fitting with constant hazards will overestimate the average hazard and select too conservative of a dose. On the other hand, if hazard rates are truly increasing, a constant hazard model will underestimate the average hazard and may choose too high of a dose.

Therefore, including a shape parameter in the hazard specification becomes necessary complexity. A Weibull-style hazard, while retaining dose on the skeleton, offers a solution. Let i represent one of a sample of n patients:

$$h_1(t_1|x_j) = \alpha\lambda_1 x_{ij} t_{1i}^{\alpha-1} \tag{2.7}$$

$$h_2(t_2|x_j) = \alpha\lambda_2 x_{ij} t_{2i}^{\alpha-1} \tag{2.8}$$

$$h_3(t_2|t_1, x_j) = \alpha\lambda_2 x_{i,j-1}^{exp(\beta_1)} t_{2i}^{\alpha-1} \tag{2.9}$$

The parameter α is a shape parameter defined on $(0, \infty)$. If $\alpha > 1$, the hazard rate is increasing over time, while $0 < \alpha < 1$ means the hazard rate is decreasing. If $\alpha = 1$, the hazards above reduce to the Constant-Skeleton specification.

Other Possible Specifications

Other specifications are possible for an illness-death model. Accelerated failure time (AFT) models are common in survival analysis. The AFT directly models sur-

vival time (log of survival time) rather than hazard rates with a linear representation.

$$log(T_{1i}) = \beta_1 x_i + \epsilon_{i1}, \quad T_{1i} > 0 \tag{2.10}$$

$$log(T_{2i}) = \beta_1 x_i + \epsilon_{i2}, \quad T_{2i} > 0 \tag{2.11}$$

$$log(T_{2i}) = \beta_1 x_i + \beta_2 \epsilon_{i2}, \quad T_{2i} > T_{1i} \tag{2.12}$$

The error distributions ϵ_{i1} and ϵ_{i2} require specification, and with at least one parameter each, the model contains at least four parameters, but with less flexibility compared to the four-parameter hazards model. The AFT and conditional hazard models are the same only if ϵ has an Extreme Value Distribution and the conditional hazards are Cox-style with Weibull (Exponential/Constant) baseline hazards.

One reason to avoid hazard specifications in general and utilizing the AFT is the difficulty in interpreting the hazard ratio commonly used to assess the effect of covariates in survival analysis. This topic has been given noticeable attention in the literature (Wei 1992, Hernán 2010, Uno et al. 2015). However, the current problem does not use the hazard ratio in any way. Baseline hazards are fully specified and event probabilities are directly calculated. Relative effects of covariates are not of interest. Additive hazard models can also be considered to estimate survival curves, but are not considered because of the same necessary linear representation mentioned above.

Dose Interruptions

One of the main differences between the current method and TITECRM-MC is the inclusion of Dose Interruptions versus Dose Reductions. Dose Interruptions are a different attempt to manage chronic low-grade toxicities, yet they still signal a relevant toxicity has occurred, and are incorporated into study protocols at least

as often as Dose Reductions. Summary statistics of these two events are typically reported together. The LURET trial reported that 16 out of 19 patients experienced a Dose Interruption while 10 experienced a Dose Reduction. In TITECRM-MC, Dose Interruptions could be included in the same outcome as a Dose Reduction. The variable $Z = 1$ then, if the patient has a Dose Interruption or a Dose Reduction. In this way, both events are treated identically and no distinction is made between the differing dosing regimes induced by the events. Conversely, an additional level of Z could be included such that $Z = 1$ for Dose Interruption, $Z = 2$ for Dose Reduction, and $Z = 3$ for DLT. This assumes Dose Reductions as distinctly worse toxic events than Dose Interruptions and requires another parameter. Because no reduced dose can be explicitly included in the TITECRM-MC model, this model is more suited for Dose Interruptions only.

With the proposed method, Dose Interruptions can be modeled more naturally. Let i represent one of a sample of n patients, and let $w_i = 1$ if patient i has a Dose Reduction, and $w_i = 0$ if patient i has a Dose Interruption:

$$h_1(t_1|x_j) = \alpha\lambda_1 x_{ij} t_{1i}^{\alpha-1} \tag{2.13}$$

$$h_2(t_2|x_j) = \alpha\lambda_2 x_{ij} t_{2i}^{\alpha-1} \tag{2.14}$$

$$h_3(t_2|t_1, x_j) = \alpha\lambda_2 x_{i,j-w_i}^{exp(\beta_1)} t_{2i}^{\alpha-1} \tag{2.15}$$

The hazard h_3 applies after a Dose Interruption or Dose Reduction, but the dose expressed $x_{i,j-w_i}$ depends on which event occurred. It is possible for a patient to have both a Dose Interruption and Dose Reduction. In that case, the five-transition model would be needed, where the first intermediate event is a Dose Interruption and the second intermediate event is a Dose Reduction. A middle ground of complexity would involve the five-transition model, with five transitions but needing only four unique

hazard rates needed (see Appendix B).

Dose Interruptions (and often Dose Reductions as well) imply a period of time when the patient does not receive the therapy, perhaps a week or two. Some clarity should be provided with regard to this period in terms of time scale. It can be assumed the patient has no chance of experiencing a DLT during this period. Thus, the interrupted period could be treated as if it never happened. A patient who is Dose Interrupted at week 14 and resumes in week 16 (real time) would have their analysis time scale resume at week 14. Alternatively, the time scale could continue to track real time. The first accommodation should be preferred, but the overall effect on dose selection would likely be minimal in any realistic scenario.

2.2 Observed Data Likelihood

The observed data likelihood of the illness-death model for an individual patient arises from four possible outcomes. Whereas for the five-transition model, this expands to six possible outcomes. The survival functions S_1, S_2, and S_3 correspond to their respectively labeled hazard functions, where $S(t) = exp\ (-\int_0^t h(u)du)$. Let C_i be the time of censoring for patient i if the patient remains on the trial at the time of analysis. The four possible outcomes are:

1. Censored before DLT or DR: $L_{1i}(\alpha, \beta_1, \lambda_1, \lambda_2) = S_1(C_i)S_2(C_i)$
2. DLT before DR: $L_{2i}(\alpha, \beta_1, \lambda_1, \lambda_2) = h_2(t_{21})S_1(C_i)S_2(C_i)$
3. Censored following first DR: $L_3 i(\alpha, \beta_1, \lambda_1, \lambda_2) = h_1(t_{1i})S_1(t_{1i})S_2(t_{1i})\frac{S_3(C_i)}{S_3(t_{1i})}$
4. DLT following first DR: $L_4 i(\alpha, \beta_1, \lambda_1, \lambda_2) = h_1(t_{1i})h_3(t_{2i}|t_{1i})S_1(t_{1i})S_2(t_{1i})\frac{S_3(t_{2i})}{S_3(t_{1i})}$

These outcomes lead to the likelihood function:

$$L(\alpha, \beta_1, \lambda_1, \lambda_2) = \prod_{i \in L_{1i}} L_{1i} \prod_{i \in L_{2i}} L_{2i} \prod_{i \in L_{3i}} L_{3i} \prod_{i \in L_{4i}} L_{4i} \tag{2.16}$$

Using L_{4i} as an example, the reasoning of each likelihood element can be explained as follows (derivations are shown in Chapter 3). The patient experiences a DLT after a DR. The DR occurs at t_{1i}, contributing $h_1(t_{1i})$. The patient must have survived a DR and a DLT up until t_{1i}, contributing $S_1(t_{1i})S_2(t_{1i})$. Then, starting at t_{1i}, the patient must have survived a DLT until t_{2i}, contributing $\frac{S_3(t_{2i})}{S_3(t_{1i})}$. The other likelihood elements can be explained in a similar fashion. If a DR occurs, the original dose is used for h_1, h_2, S_1, and S_2, while the reduced dose is used for h_3 and S_3.

2.3 Prior Distributions

Under a Bayesian framework, a prior distribution must be specified for λ_1, λ_2, and β_1 in the Constant-Skeleton specification, with an additional prior for α in the Weibull-Skeleton. A simple, uninformative choice is to specify a uniform prior for each parameter, with an upper bound on each of α, λ_1, and λ_2 that is above all plausible values, and a lower and upper bound on β_1.

$$\alpha \sim \text{Uniform}(0, b_\alpha) \tag{2.17}$$

$$\lambda_1 \sim \text{Uniform}(0, b_{\lambda_1}) \tag{2.18}$$

$$\lambda_2 \sim \text{Uniform}(0, \lambda_1) \tag{2.19}$$

$$\beta_1 \sim \text{Uniform}(a_{\beta_1}, b_{\beta_1}) \tag{2.20}$$

A second choice for priors that can remain relatively uninformative, but provide more realistic values for parameter when little data is available is a truncated Normal,

$$\alpha \sim N_{+}(\mu_{\alpha}, \sigma_{\alpha}^{2}) \tag{2.21}$$

$$\lambda_1 \sim N_{+}(\mu_{\lambda_1}, \sigma_{\lambda_1}^{2}) \tag{2.22}$$

$$\lambda_2 \sim N_{+,\lambda_1}(\mu_{\lambda_2}, \sigma_{\lambda_2}^{2}) \tag{2.23}$$

$$\beta_1 \sim N(\mu_{\beta_1}, \sigma_{\beta_1}^{2}), \tag{2.24}$$

where α, λ_1, and λ_2 are bounded positive and λ_2 is also bounded above by λ_1. The suggestion of these priors is partially motivated by the unique formulation of the hazards. There is no experiential knowledge to lean on in choosing more specialized prior distributions that may be more cohesive with the model. Moreover, this situation represents a real-life scenario where a study team is likely to choose the most basic, common prior distributions for fear of choosing overly specific or complex prior distributions. The simulations in Chapter 4 cannot then be accused of using bespoke prior distributions nor of benefiting in performance versus comparators.

Given data from j patients on the trial and the likelihood, the joint posterior distribution of the parameters can be given in the general form using Bayes' Rule,

$$p(\tau|D_j) = \frac{L(\tau|D_j)h(\tau)}{p(D_j)} \tag{2.25}$$

$$p(\tau|D_j) = \frac{L(\tau|D_j)h(\tau)}{\int L(\tau|D_j)h(\tau)d\tau}, \tag{2.26}$$

where $\tau = \{\lambda_1, \lambda_2, \beta_1\}$ or $\tau = \{\alpha, \lambda_1, \lambda_2, \beta_1\}$. The Maximum a posteriori (MAP) estimate from the posterior distribution of τ can be used to obtain $\hat{\tau}$. The MAP estimate for a parameter maximizes the probability of the posterior distribution of τ,

or equivalently stated, finds the mode of the posterior distribution,

$$\hat{\tau} = argmax_{\tau} P(\tau|D_j). \tag{2.27}$$

2.4 Transition Probabilities

To answer clinical questions of interest, the estimated conditional hazard and survival functions can be used to estimate the conditional probability of a future state given a current state, or rather transition probabilities. The two most relevant transition probabilities in a clinical setting might be P(Healthy→DR) and P(Healthy→DLT), via any possible path. The probabilities are specified below, following the form of Putter, Fiocco, and Geskus 2007. Transition probabilities for the five-transition model with explicit forms for the Constant-Skeleton and Weibull-Skeleton specifications, are in Appendix B. Let t_c be the time at which patients begin observation, likely $t_c = 0$, and let t be the future time of interest in the prediction.

1. P(Healthy → DR)

$$\begin{aligned} P(H \to DR) &= P_{t_c}(T_1 \leq t, T_2 > T_1 | T_1 > t_c, T_2 > t_c) \\ &= \frac{\int_{t_c}^{t} h_1(r) S_1(r) S_2(r) dr}{S_1(t_c) S_2(t_c)} \end{aligned} \tag{2.28}$$

2. P(Healthy → DLT) directly

$$\begin{aligned} P(H \to DLT) &= P_{t_c}(T_1 > T_2, T_2 \leq t | T_1 > t_c, T_2 > t_c) \\ &= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r) dr}{S_1(t_c) S_2(t_c)} \end{aligned} \tag{2.29}$$

3. P(Healthy → DLT) via any possible path

$$
\begin{aligned}
&P(H \to DLT) + P(H \to DR \to DLT) \\
&= P_{t_c}(T_2 \le t | T_1 > t_c, T_2 > t_c) \\
&= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r) dr}{S_1(t_c) S_2(t_c)} + \frac{\int_{t_c}^{t} h_1(r) S_1(r) S_2(r) \frac{\int_r^t h_3(u) S_3(u) du}{S_3(r)} dr}{S_1(t_c) S_2(t_c)}
\end{aligned} \tag{2.30}
$$

Using transition probability P(Healthy→DLT) directly as an example, a reasoning of the transition probability expressions is similar to that of the likelihood scenarios. A patient, starting in a Healthy state at time t_c, must survive a DR and a DLT up until time r, sometime between t_c and t, contributing $S_1(r)S_2(r)$. At that point, the patient experiences a DR, contributing $h_1(r)$. From the time of DR, the patient must then survive a DLT until time t, contributing $\frac{S_3(t)}{S_3(r)}$. The probability P(H → DR) is importantly expressed as the probability of having a Dose Reduction by time t, not the probability of being in a Dose Reduced state by time t. This means a patient could also have a DLT by time t, and would still be counted as having a DR. The quantity of interest is how likely it is for a patient to make the transition, not finish in the state. This is also where the current method fundamentally differs from the TITECRM-MC in prediction.

2.5 Dose Assignment and Selection

To adaptively assign a dose to new patients entering a trial, as well as ultimately select one or more doses at the end of a trial, various dose assignment and selection rules could be applied using estimates from the explained model. A simple approach

is minimal "Total Distance" estimated from each dose to two pre-specified toxicity targets. Let π_{DR} and π_{DLT} be toxicity targets for each respective event, between 0 and 1. The dose estimated to have the smallest total distance is defined as

$$argmax_j(|\hat{\pi}_{DR}(x_j) - \pi_{DR}| + |\hat{\pi}_{DLT}(x_j) - \pi_{DLT}|), \tag{2.31}$$

where $\hat{\pi}_{DR}(x_j)$ is the estimated probability of Dose Reduction by end of follow-up at dose x_j; and $\hat{\pi}_{DLT}(x_j)$ is the estimated probability of DLT. This dose would be assigned to the next patient brought onto the trial, or if at the end of the trial, chosen as the dose to move forward. If multiple forward doses are desired k, the doses with the k smallest estimated distances would be chosen. This approach does not penalize doses for having estimated probabilities above the targets. The chosen dose may be above both targets, below both targets, or above one and below the other.

A modification of the Total Distance approach is Weighted Total Distance. A linear weight ω is applied to the DLT distance

$$argmin_j(|\hat{\pi}_{DR}(x_j) - \pi_{DR}| + \omega|\hat{\pi}_{DLT}(x_j) - \pi_{DLT}|), \tag{2.32}$$

so that the DLT distance becomes more or less important, relative to the DR distance in choosing the best dose. Most often under this approach, ω would be greater than 1 to make DLT distance relatively more important. (Note that until now, DLTs have been treated with exclusive importance.) A more conservative approach, one that treats π_{DR} and π_{DLT} as upper bounds or thresholds instead of targets, is choosing the maximum dose with an estimated probability of DLT under a target θ. Call this

the Maximum Uniformly Tolerated Dose (MUTD) x_j^*,

$$x_j^* = max\{x_j \in X : \hat{\pi}_{DLT}(x_j) \leq \pi_{DLT}, \hat{\pi}_{DR}(x_j) \leq \pi_{DR}\}. \tag{2.33}$$

The MUTD, like many dose-finding designs and CRM extensions, is cautious of overdosing. The dose remains low throughout a trial, leading to difficulty exploring the high end of the dose range, especially when sample sizes are small. With targeted agents providing a different toxicity profile, it may be more beneficial to relax control on overdosing in order to better explore the higher end of the dose range. Therefore, a combination strategy of Total Distance and MUTD may be the most prudent, as π_{DR} and π_{DLT} are still important thresholds. During a trial, each new patient is dosed based on Total Distance, with a weight if desired. At the end of the trial, a dose or multiple doses are selected to move forward based on the MUTD.

Often with dose-finding models, an adaptive randomization stage is used at the beginning of the trial to dose patients. Researchers are cautious that there may not be enough data early on to trust the estimates of a dose-finding model. An adaptive randomization stage is a separate algorithm that encourages controlled exploration of the dose range. The sample size used in the adaptive randomization stage is based on preference of the researchers. No adaptive randomization stage is proposed for this method to reduce complexity of the algorithm in total, primarily because a suitable prior and dose skeleton, which must be specified anyway, can go lengths in encouraging or discouraging dose exploration.

Chapter 3

Theoretical Results

Considering semi-competing risks data, the joint distribution of (T_1, T_2), defined in this case by conditional hazard functions, is fully identified only when $(0 < t_1 < t_2)$. Consequently, $P(T_1 < \infty) \leq 1$ (Lee et al. 2015). A resolution is to set $T_1 = \infty$ if the terminal event T_2 occurs before T_1. Xu, Kalbfleisch, and Tai 2010 argue this strategy is preferable to assuming a latent distribution of (T_1, T_2) over the region $(t_1 > t_2)$, and a more true reflection of the physical situation.

3.1 Likelihood and Joint Density

The observed data likelihood for the three-transition illness-death model can be derived generally via joint density for T_1 and T_2. Thinking about the paths an observed data point can take, the joint density can be separated where T_2 comes before T_1 and where T_2 comes after T_1. Let t_1 and t_2 be the observed times of toxicity.

$$
\begin{aligned}
&P(T_1 = t_1, T_2 = t_2) \\
&= P(T_1 = t_1, T_2 = t_2, T_2 > t_1) + && (3.1) \\
&\quad P(T_1 = t_1, T_2 = t_2, T_1 > t_2) && (3.2) \\
&= f_a(t_1, t_2) + f_b(t_1, t_2) && (3.3)
\end{aligned}
$$

The two likelihood elements f_a and f_b can then be derived separately. Let h_1, h_2, and h_3 be defined as hazard functions of Dose Reduction, DLT before Dose Reduction, and DLT after Dose Reduction. Let S_1, S_2, and S_3 be the associated survival functions.

$$h_1(t_1) = \frac{P(T_1 = t_1 | T_2 > t_1)}{P(T_1 > t_1 | T_2 > t_1)} \tag{3.4}$$

$$h_2(t_2) = \frac{P(T_2 = t_2 | T_1 > t_2)}{P(T_2 > t_2 | T_1 > t_2)} \tag{3.5}$$

$$h_3(t_2|t_1) = \frac{P(T_2 = t_2 | T_1 = t_1, T_2 > t_1)}{P(T_2 > t_2 | T_1 = t_1, T_2 > t_1)} \tag{3.6}$$

$$S_1(t_1) = P(T_1 > t_1) \tag{3.7}$$

$$S_2(t_2) = P(T_2 > t_2) \tag{3.8}$$

$$S_3(t_2|t_1) = P(T_2 > t_2 | T_1 = t_1) \tag{3.9}$$

Joint Density Piece 1. $f_a(t_1, t_2) = h_1(t_1) S_1(t_1) S_2(t_1) h_3(t_2|t_1) \dfrac{S_3(t_2|t_1)}{S_3(t_1|t_1)}$.

Derivation:

$P(T_1 = t_1, T_2 = t_2, T_2 > t_1)$

$$= P(T_1 = t_1, T_2 > t_1) P(T_2 = t_2 | T_1 = t_1, T_2 > t_1) \tag{3.10}$$

$$= P(T_1 = t_1, T_2 > t_1) \frac{P(T_1 > t_1, T_2 > t_1)}{P(T_1 > t_1, T_2 > t_1)} P(T_2 = t_2 | T_1 = t_1, T_2 > t_1) \tag{3.11}$$

$$= \frac{P(T_1 = t_1, T_2 > t_1)}{P(T_1 > t_1, T_2 > t_1)} P(T_1 > t_1, T_2 > t_1) \frac{P(T_2 = t_2 | T_1 = t_1, T_2 > t_1)}{P(T_2 > t_2, T_1 = t_1, T_2 > t_1)} \times \tag{3.12}$$

$$P(T_2 > t_2, T_1 = t_1, T_2 > t_1)$$

$$= \frac{P(T_1 = t_1, T_2 > t_1)}{P(T_1 > t_1, T_2 > t_1)} P(T_1 > t_1, T_2 > t_1) \frac{P(T_2 = t_2 | T_1 = t_1, T_2 > t_1)}{P(T_2 > t_2, T_1 = t_1, T_2 > t_1)} \times \tag{3.13}$$

$$\frac{P(T_2 > t_2, T_2 > t_1 | T_1 = t_1)}{P(T_2 > t_1 | T_1 = t_1)}$$

In the expression $P(T_2 > t_2, T_2 > t_1 | T_1 = t_1)$, because t_1 is given and finite, it must be allowed to happen, thus $T_2 > t_1$. The expression can then be simplified to $P(T_2 > t_2 | T_1 = t_1)$,

$$= \frac{P(T_1 = t_1, T_2 > t_1)}{P(T_1 > t_1, T_2 > t_1)} P(T_1 > t_1, T_2 > t_1) \times \frac{P(T_2 = t_2 | T_1 = t_1, T_2 > t_1)}{P(T_2 > t_2, T_1 = t_1, T_2 > t_1)} \frac{P(T_2 > t_2 | T_1 = t_1)}{P(T_2 > t_1 | T_1 = t_1)} \tag{3.14}$$

$$= h_1(t_1) S_1(t_1) S_2(t_1) h_3(t_2 | t_1) \frac{S_3(t_2 | t_1)}{S_3(t_1 | t_1)} \tag{3.15}$$

Note, this derivation is based on the conditional independence of the joint survival function $P(T_1 > t_1, T_2 > t_2)$, given a frailty parameter γ to account for the dependence between T_1 and T_2 (Rodriguez 2010). Thus, a rigorous application of this method would include a frailty parameter. However, for reasons stated previously in Section 2.1, the current application of this method foregoes a frailty parameter for simplicity and efficiency.

Joint Density Piece 2. $f_b(t_1, t_2) = h_2(t_2) S_1(t_2) S_2(t_2)$

Derivation:

$P(T_1 = t_1, T_2 = t_2, T_1 > t_2)$

$$= P(T_1 = t_1 | T_2 = t_2, T_1 > t_2) P(T_2 = t_2, T_1 > t_2) \tag{3.16}$$

$$= P(T_1 = t_1 | T_2 = t_2, T_1 > t_2) \frac{P(T_2 = t_2, T_1 > t_2)}{P(T_2 > t_2, T_1 > t_2)} P(T_2 > t_2, T_1 > t_2) \tag{3.17}$$

If $T_1 > t_2$, then t_1 is assumed to be ∞.

$$= P(T_1 = \infty | T_2 = t_2, T_1 > t_2) \frac{P(T_2 = t_2, T_1 > t_2)}{P(T_2 > t_2, T_1 > t_2} P(T_2 > t_2, T_1 > t_2) \tag{3.18}$$

$$= 1 * h_2(t_2) S_2(t_2) S_1(t_2) \tag{3.19}$$

These two likelihood elements can then be used to specify the four unique paths of the illness-death model. The four paths are: 1) Censored before any toxicity. 2) DLT before Dose Reduction. 3) Censored after Dose Reduction. 4) DLT after Dose Reduction. Let C_i represent a time of censoring. Each path is expressed in terms of hazard and survival functions, thus (2.4), (2.5), and (2.6) or (2.7), (2.8), and (2.9) can be directly substituted with their associated survival functions to give specifications for the current dose-finding context. These explicit forms are given in Appendix B.

1. Censored before any toxicity: $L_{1i} = P(T_1 > C_i, T_2 > C_i) = S_1(C_i) S_2(C_i)$

Derivation: Direct result from joint survivor function as explained in (3.14).

2. DLT before Dose Reduction: $L_{2i} = P(T_1 = \infty, T_2 = t_{2i}, T_1 > t_{2i})$

$$= h_2(t_{2i}) S_1(t_{2i}) S_2(t_{2i})$$

Derivation:

$P(T_1 = \infty, T_2 = t_{2i}, T_1 > t_{2i})$

$$= P(T_1 = t_1, T_2 = t_2, T_1 > t_2) \tag{3.20}$$

$$= f_b(t_1, t_2) \tag{3.21}$$

$$= h_2(t_2) S_1(t_2) S_2(t_2) \tag{3.22}$$

3. Censored after Dose Reduction: $L_3 i = h_1(t_{1i})S_1(t_{1i})S_2(t_{1i})\dfrac{S_3(C_i|t_{1i})}{S_3(t_{1i}|t_{1i})}$

Derivation:

$$\int_{C_i}^{\infty} f_a(t_{1i}, s)ds$$

$$= h_1(t_{1i})S_1(t_{1i})\frac{S_2(t_{1i})}{S_3(t_{1i})}\int_{C_i}^{\infty} h_3(s|t_{1i})S_3(s)ds \tag{3.23}$$

$$= h_1(t_{1i})S_1(t_{1i})\frac{S_2(t_{1i})}{S_3(t_{1i})}\int_{C_i}^{\infty} h_3(s|t_{1i})exp(-\int_0^u h_3(u|t_{1i}))duds \tag{3.24}$$

$$= h_1(t_{1i})S_1(t_{1i})\frac{S_2(t_{1i})}{S_3(t_{1i})}[-exp(-\int_0^u h_3(u|t_{1i})du|_0^{C_i}] \tag{3.25}$$

$$= h_1(t_{1i})S_1(t_{1i})\frac{S_2(t_{1i})}{S_3(t_{1i})}[-0 + S_3(C_i)] \tag{3.26}$$

$$= h_1(t_{1i})S_1(t_{1i})S_2(t_{1i})\frac{S_3(C_i)}{S_3(t_{1i})} \tag{3.27}$$

4. DLT after Dose Reduction: $L_4 i = h_1(t_{1i})h_3(t_{2i}|t_{1i})S_1(t_{1i})S_2(t_{1i})\dfrac{S_3(t_{2i}|t_{1i})}{S_3(t_{1i}|t_{1i})}$

Derivation:

$$P(T_1 = t_{1i}, T_2 = t_{2i}, T_2 > t_{1i})$$

$$= f_a(t_1, t_2) \tag{3.28}$$

$$= h_1(t_{1i})S_1(t_{t1i})S_2(t_{1i})h_3(t_{2i}|t_{1i})\frac{S_3(t_{2i})}{S_3(t_{1i})} \tag{3.29}$$

3.2 Transition Probabilities

The transition probabilities are derived in a similarly straightforward manner. Let t be the transition time of interest, and let t_c be the time at which patients started

observation.

Healthy $\rightarrow$ DR: $P_{t_c}(T_1 \leq t, T_2 > T_1 | T_1 > t_c, T_2 > t_c) = \dfrac{\int_{t_c}^{t} h_1(r) S_1(r) S_2(r) dr}{S_1(t_c) S_2(t_c)}$

Derivation:

$$
\begin{aligned}
& P_{t_c}(T_1 \leq t, T_2 > T_1 | T_1 > t_c, T_2 > t_c) \\
&= \frac{P_{t_c}(T_1 \leq t, T_2 > T_1, T_1 > t_c, T_2 > t_c)}{P_{t_c}(T_1 > t_c, T_2 > t_c)} && (3.30) \\
&= \frac{P_{t_c}(t_c < T_1 \leq t, T_2 > T_1)}{S_1(t_c) S_2(t_c)} && (3.31) \\
&= \frac{\int_{t_c}^{t} P(T_1 = r, T_2 > r) dr}{S_1(t_c) S_2(t_c)} && (3.32) \\
&= \frac{\int_{t_c}^{t} P(T_1 = r, T_2 > r) \dfrac{P(T_1 > r, T_2 > r)}{P(T_1 > r, T_2 > r)} dr}{S_1(t_c) S_2(t_c)} && (3.33) \\
&= \frac{\int_{t_c}^{t} h_1(r) P(T_1 > r, T_2 > r) dr}{S_1(t_c) S_2(t_c)} && (3.34) \\
&= \frac{\int_{t_c}^{t} h_1(r) S_1(r) S_2(r) dr}{S_1(t_c) S_2(t_c)} && (3.35)
\end{aligned}
$$

Healthy → DLT, directly: $P_{t_c}(T_1 > T_2, T_2 \leq t | T_1 > t_c, T_2 > t_c)$

$$= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r) dr}{S_1(t_c) S_2(t_c)}$$

Derivation:

$$P_{t_c}(T_2 \leq t, T_1 > T_2 | T_1 > t_c, T_2 > t_c)$$

$$= \frac{P_{t_c}(T_1 \leq t, T_2 > T_1, T_1 > t_c, T_2 > t_c)}{P_{t_c}(T_1 > t_c, T_2 > t_c)} \tag{3.36}$$

$$= \frac{P_{t_c}(t_c < T_2 \leq t, T_1 > T_2)}{S_1(t_c) S_2(t_c)} \tag{3.37}$$

$$= \frac{\int_{t_c}^{t} P(T_2 = r, T_1 > r) dr}{S_1(t_c) S_2(t_c)} \tag{3.38}$$

$$= \frac{\int_{t_c}^{t} P(T_2 = r, T_1 > r) \frac{P(T_2 > r, T_1 > r)}{P(T_2 > r, T_1 > r)} dr}{S_1(t_c) S_2(t_c)} \tag{3.39}$$

$$= \frac{\int_{t_c}^{t} h_2(r) P(T_2 > r, T_1 > r) dr}{S_1(t_c) S_2(t_c)} \tag{3.40}$$

$$= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r) dr}{S_1(t_c) S_2(t_c)} \tag{3.41}$$

Healthy → DLT via any path: $P_{t_c}(T_2 \leq t | T_1 > t_c, T_2 > t_c)$

$$= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r) dr}{S_1(t_c) S_2(t_c)} + \frac{\int_{t_c}^{t} h_1(r) S_1(r) S_2(r) \frac{\int_r^t h_3(u) S_3(u|r) du}{S_3(r|r)} dr}{S_1(t_c) S_2(t_c)}$$

Derivation:

$P_{t_c}(T_2 \leq t | T_1 > t_c, T_2 > t_c)$

$$= P_{t_c}(T_1 > T_2, T_2 \leq t | T_1 > t_c, T_2 > t_c) + P_{t_c}(T_1 < T_2, T_2 \leq t | T_1 > t_c, T_2 > t_c)$$

$$= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r)}{S_1(t_c) S_2(t_c)} + P_{t_c}(T_1 < T_2, T_2 \leq t | T_1 > t_c, T_2 > t_c) \tag{3.42}$$

$$= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r)}{S_1(t_c) S_2(t_c)} + \frac{P_{t_c}(T_1 < T_2, T_2 \leq t, T_1 > t_c, T_2 > t_c)}{P_{t_c}(T_1 > t_c, T_2 > t_c)} \tag{3.43}$$

$$= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r)}{S_1(t_c) S_2(t_c)} + \frac{P_{t_c}(t_c < T_1 < t, T_1 < T_2 \leq t)}{P_{t_c}(T_1 > t_c, T_2 > t_c)} \tag{3.44}$$

$$= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r)}{S_1(t_c) S_2(t_c)} + \frac{\int_{t_c}^{t} P(T_1 = r, r < T_2 < t) dr}{S_1(t_c) S_2(t_c)} \tag{3.45}$$

$$= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r)}{S_1(t_c) S_2(t_c)} + \frac{\int_{t_c}^{t} P(T_1 = r, r < T_2 < t) \frac{P(T_1 > r, r < T_2 < t)}{P(T_1 > r, r < T_2 < t)} dr}{S_1(t_c) S_2(t_c)} \tag{3.46}$$

$$= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r)}{S_1(t_c) S_2(t_c)} + \frac{\int_{t_c}^{t} h_1(r) P(T_1 > r, r < T_2 < t) dr}{S_1(t_c) S_2(t_c)} \tag{3.47}$$

$$= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r)}{S_1(t_c) S_2(t_c)} + \frac{\int_{t_c}^{t} h_1(r) P(T_1 > r, T_2 > r) P(T_2 < t | T_1 > r, T_2 > r) dr}{S_1(t_c) S_2(t_c)} \tag{3.48}$$

$$= \frac{\int_{t_c}^{t} h_2(r) S_1(r) S_2(r)}{S_1(t_c) S_2(t_c)} + \frac{\int_{t_c}^{t} h_1(r) S_1(r) S_2(r) \frac{\int_r^t h_3(u) S_3(u|r)}{S_3(r|r)} dr}{S_1(t_c) S_2(t_c)} \tag{3.49}$$

DR $\rightarrow$ DLT: $P_{t_c}(T_2 < t | T_1 = t_1 < t_c, T_2 > t_c) = \dfrac{\int_{t_c}^{t} h_3(r) S_3(r|t_c) dr}{S_3(t_c|t_c)}$

Derivation:

$$
\begin{aligned}
& P_{t_c}(T_2 \leq t | T_1 < t_c, T_2 > t_c) \\
&= \frac{P_{t_c}(T_2 \leq t, T_2 > t_c | T_1 < t_c)}{P_{t_c}(T_2 > t_c | T_1 < t_c)} && (3.50) \\
&= \frac{\int_{t_c}^{t} P(T_2 = r | T_1 < t_c) dr}{S_3(t_c|t_c)} && (3.51) \\
&= \frac{\int_{t_c}^{t} P(T_2 = r | T_1 < t_c) \dfrac{P(T_2 > r | T_1 < t_c)}{P(T_2 > r | T_1 < t_c)} dr}{S_3(t_c|t_c)} && (3.52) \\
&= \frac{\int_{t_c}^{t} h_3(r) S_3(r|t_c)}{S_3(t_c|t_c)} && (3.53)
\end{aligned}
$$

3.3 Connection to Counting Processes

As a multi-state model, the semi-competing risks framework can also be described using the language of Counting or Stochastic Processes. A stochastic process X_t, $t \in [0, \infty)$ is defined as $X_t = \ell$, if the process is in state ℓ at time t. Let F_t be a σ-algebra covering the information over $[0,\, t]$. The hazard function, alternatively named an intensity with represenation λ, of the transition from state m to state ℓ is defined as

$$\lambda_{m\ell}(t | X_u, u \in [0, t)) = \lim_{\Delta t \downarrow 0} \frac{Pr\{X_{t+\Delta t} = \ell | F_{t^-}, X_{t^-} = m\}}{\Delta t}. \tag{3.54}$$

Again, transition probabilities are typically the quantity of interest, and defining hazards or intensities are an efficient tool in calculating these probabilities. Transition

probabilities evaluated at time t can be generally defined as

$$P_\ell(\nu, t) = Pr(X_t = \ell | X_u, u \in [0, \nu]) \tag{3.55}$$

which are conditional on X_u, the whole process history up to time ν.

The ideas of Markov and Semi-Markov exist in the same way (Semi-Markov also being referred to as Markov Extension, Hougaard 1999). A Markov model considers X_u to contain information only about the current state at time ν. Semi-Markov additionally considers the time at which the process moved into the current state. Thus, in the Markov setting, the intensity and transition probabilities can be simplified as

$$\lambda_{m\ell}(t) = \lim_{\Delta t \downarrow 0} \frac{Pr\{X_{t+\Delta t} = \ell | X_{t^-} = m\}}{\Delta t} \tag{3.56}$$

$$P_{m\ell}(\nu, t) = Pr(X_t = \ell | X_\nu = m) \tag{3.57}$$

The total hazard of leaving state m is $\lambda_m(t) = \sum_{\ell \neq m} \lambda_{ml}(t)$. The likelihood for the Markov model comes from Hougaard 1999. Let an observation have E events from time 0 to C experienced at times $j = \{1, ..., E\}$. The process starts at state s_0 and moves into states s_j. The likelihood of a single observation is defined as

$$\left[\prod_{j=1}^{E} \lambda_{s_{j-1}s_j}(t_j) exp\{-\int_{t_{j-1}}^{t_j} \lambda_{s_{j-1}}(u)du\}\right] exp\{-\int_{t_E}^{C} \lambda_{S_E}(u)du\}. \tag{3.58}$$

This specification has the same explanation as the likelihood previously. Transitions are experienced at times $j = \{1, ..., E\}$, represented by the intensities. Then, within each state, all other possible hazards leading out of that state are "survived" until the next transition time, represented by the exponential terms (survival functions). The likelihood for each of the four paths in an illness-death model is formulated below

using this notation.

1. Censored before any toxicity: E = 0

$$L_1 = \left[\prod^{0} \lambda_{s_{j-1}s_j}(t_j) exp\{-\int_{t_{j-1}}^{t_j} \lambda_{s_{j-1}}(u)du\}\right] exp\{-\int_0^C (\lambda_1(u)+\lambda_2(u))du\} \quad (3.59)$$

$$= exp\{-\int_0^C \lambda_1(u)du\}\ exp\{-\int_0^C \lambda_2(u)du\} \quad (3.60)$$

2. DLT before Dose Reduction: E = 1

$$L_2 = \left[\prod^{1} \lambda_{s_{j-1}s_j}(t_j) exp\{-\int_{t_{j-1}}^{t_j} \lambda_{s_{j-1}}(u)du\}\right] \quad (3.61)$$

$$= \lambda_2(t_2)\ exp\{-\int_0^C (\lambda_1(u)+\lambda_2(u))du\} \quad (3.62)$$

$$= \lambda_2(t_2)\ exp\{-\int_0^C \lambda_1(u)du\}\ exp\{-\int_0^C \lambda_2(u)du\} \quad (3.63)$$

3. Censored after Dose Reduction: E = 1

$$L_3 = \left[\prod^{1} \lambda_{s_{j-1}s_j}(t_j) exp\{-\int_{t_{j-1}}^{t_j} \lambda_{s_{j-1}}(u)du\}\right] exp\{-\int_{t_1}^C \lambda_3(u)du\} \quad (3.64)$$

$$= \lambda_1(t_1)\ exp\{-\int_0^{t_1} (\lambda_1(u)+\lambda_2(u))du\} exp\{-\int_{t_1}^C \lambda_3(u)du\} \quad (3.65)$$

$$= \lambda_2(t_2)\ exp\{-\int_0^{t_1} \lambda_1(u)du\}\ exp\{-\int_0^{t_1} \lambda_2(u)du\}\ exp\{-\int_{t_1}^C \lambda_3(u)du\} \quad (3.66)$$

4. DLT after Dose Reduction: E = 2

$$L_4 = \left[\prod_{j=1}^{2} \lambda_{s_{j-1}s_j}(t_j) exp\{-\int_{t_{j-1}}^{t_j} \lambda_{s_{j-1}}(u)du\}\right] \tag{3.67}$$

$$= \left[\lambda_1(t_1)\ exp\{-\int_0^{t_1} (\lambda_1(u) + \lambda_2(u))du\}\right] \left[\lambda_3(t_2) exp\{-\int_{t_1}^{t_2} \lambda_3(u)du\}\right] \tag{3.68}$$

$$= \lambda_1(t_1)\lambda_3(t_2)\ exp\{-\int_0^{t_1} \lambda_1(u)du\}\ exp\{-\int_0^{t_1} \lambda_2(u)du\}\ exp\{-\int_{t_1}^{t_2} \lambda_3(u)du\} \tag{3.69}$$

Thus, each potential path and the resulting full likelihood is confirmed to be identical as that derived previously. As for asymptotics, in applying these models, one often does not consider whether the asymptotic distributions really apply (Hougaard 1999). Especially in the current setting, sample sizes are unavoidably small, and large sample properties are not relevant. Additionally, nice asymptotic distributions are not always possible with multi-state models. For completeness, though, the next section derives an asymptotic result of choosing the correct dose based on the current model.

3.4 Asymptotic Properties

A brief result is shown to guarantee the selection of the correct dose with infinite data, including discussion of an alternative version of the same result.

Result

The observed data are assumed to follow an illness-death model defined by conditional hazard functions, independent and identically distributed between subjects.

Let there be J doses under consideration x_j, $j \in \{1, ...J\}$. The probability of a patient experiencing a DR or DLT by a time of interest t_{max} is conditional upon initial dose and expressed previously by (3.33) and (3.47).

$$p_{(DR,\ j)} = \int_0^{t_{max}} h_1(r|x_j)S_1(r|x_j)S_2(r|x_j)dr \tag{3.70}$$

$$p_{(DLT,\ j)} = \int_0^{t_{max}} h_1(r|x_j)S_1(r|x_j)S_2(r|x_j)dr + \int_0^{t_{max}} h_1(r|x_j)S_1(r|x_j)S_2(r|x_j)\frac{\int_r^{t_{max}} h_3(u|x_{j-1})S_3(u|x_{j-1})du}{S_3(r|x_{j-1})}dr \tag{3.71}$$

There exists a desired dose x_j^* in the set X such that,

$$x_j^* = max\{x_j \in X :\ p_{(DR,\ j)} \leq \pi_{DR},\ \ p_{(DLT,\ j)} \leq \pi_{DLT}\}, \tag{3.72}$$

where π_{DR} and π_{DLT} are prespecified probability targets between 0 and 1.

The proof can proceed in two steps. First, the posterior distribution of the model converges to the true parameter values as shown by Gelman et al. 2013. Then, the estimated probabilities of each event $\hat{p}_{(DR,\ j)}$, $\hat{p}_{(DLT,\ j)}$, plugging in MAP estimates, converges to the true probabilities via the continuous mapping theorem, whereby the correct dose will be selected.

Step 1: Let there be a vector of parameters θ defined on a continuous space Θ for which there is a prior distribution $p(\theta)$ and a likelihood $p(y|\theta)$. Assuming the likelihood model is correct, there is some true parameter value θ_0 in the space.

Define a neighborhood of θ_0 as an open set of all points in Θ within a fixed nonzero distance of θ_0. Let A be a neighborhood of θ_0 with nonzero prior probability. Place a neighborhood around each point in Θ, e.g. (A, B_1, B_2,...B_n), such that θ_0 is exclusively in A. Let Θ be compact, so that a finite subset of neighborhoods covering

Θ can always be chosen.

First, KL divergence at any value θ is defined as

$$KL(\theta) = E(log(\frac{f(y_i)}{p(y_i|\theta)})) = \int log(\frac{f(y_i)}{p(y_i|\theta)}) f(y_i) dy_i. \tag{3.73}$$

where f is the true distribution of the data and p is the probability model. By Jensen's Inequality,

$$\begin{aligned} \int log(\frac{f(y_i)}{p(y_i|\theta)}) f(y_i) dy_i &= \int -log(\frac{p(y_i|\theta)}{f(y_i)}) f(y_i) dy_i && (3.74) \\ &\geq -log(\int \frac{p(y_i|\theta)}{f(y_i)} f(y_i) dy_i) && (3.75) \\ &= -log(\int p(y_i|\theta) dy_i) && (3.76) \\ &= -log(1) && (3.77) \end{aligned}$$

Thus, KL divergence is greater than or equal to zero. If $\theta = \theta_0$, then $\frac{f(y_i)}{p(y_i|\theta)} = 1$ and KL divergence becomes

$$\begin{aligned} \int log(1) f(y_i) dy_i \geq log(\int f(y_i) dy_i) && (3.78) \\ \int 0 dy_i = 0 \geq log(1) = 0. && (3.79) \end{aligned}$$

and the true parameter value θ_0 necessarily minimizes the KL divergence as $KL(\theta) = 0$.

A posterior log odds-type expression relative to A is

$$log(\frac{p(\theta|y)}{p(\theta \in A|y)}) = log(\frac{p(\theta)}{p(\theta \in A)}) + \sum_{i=1}^{n} log(\frac{p(y_i|\theta)}{p(y_i|\theta \in A)}). \tag{3.80}$$

Taking the expectation of each term in the summation gives

$$E(log(\frac{p(y_i|\theta)}{p(y_i|\theta \in A)})) = KL(\theta \in A) - KL(\theta). \tag{3.81}$$

This expression is zero if $\theta \in A$ and negative in all other scenarios as θ_0 is the minimizing value of $KL(\theta)$. If $\theta \notin A$, then the summation is of random variables with negative mean. By Law of Large Numbers, the summation goes to $-\infty$ as $n \to \infty$.

If $P(\theta \in A) > 0$, then the term $log(\frac{p(\theta)}{p(\theta \in A)})$ is finite, and the right side of (3.89) goes to $-\infty$. Thus, $\frac{p(\theta)}{p(\theta \in A)}$ on the left side $\to 0$, and $p(\theta) \to 0$. Then it must be that $P(\theta \in A) \to 1$. The posterior converges to the true parameters with infinite sample size.

Step 2: Let the vector of parameters of the stated illness-death model be defined on a continuous space. With infinite data, assuming the likelihood model is correct, the vector of parameters will converge to the true parameters. The expressions $\hat{p}_{(DR,\ j)}$ and $\hat{p}_{(DLT,\ j)}$ are continuous functions of parameter estimates and known data. Thus, by the Continuous Mapping Theorem, as $\hat{\theta} \to \theta_0$, therefore $\hat{p}_{(DR,\ j)} \to p_{(DR,\ j)}$ and $\hat{p}_{(DLT,\ j)} \to p_{(DLT,\ j)}$ as $n \to \infty$. The desired dose x_j^* is guaranteed to be selected.

Alternative version

An alternative version of the same proof can be shown using only observed events, foregoing a prediction model entirely. Let there be n patients who have at least started the trial. Let the event data of these patients be represented by the set D, where the pair $D_i = \{D_{1i}, D_{2i}\}$ represents whether the patient i has experienced a DR or DLT, respectively. The number of patients n can be separated into those on each dose who have completed follow-up to t_{max} and those who have not, $n = \sum^j n_j$,

$n_j = n_{max,j} + n_{part,j}$. Assume patients can be accrued only at integer time points; also, assume a maximum of m patients can be accrued at each time point. Thus, n_{part} must be finite.

The data D_i^{max} from completed follow-up are guaranteed to each follow a Bernoulli distribution with associated probability.

$$D_{1i,\ j}^{max} \sim bern(p_{(DR,j)}) \tag{3.82}$$

$$D_{2i,\ j}^{max} \sim bern(p_{(DLT,j)}) \tag{3.83}$$

The data D_i^{part} from patients still on the trial are Bernoulli distributed, but with probability less than $p_{(DR,j)}$ or $p_{(DR,j)}$, each determined individually by the amount of time on the trial. Let X_{1j} be the number of DRs observed on dose j and X_2 be the number of DLTs observed on dose j.

$$X_{1j} = \sum^{N} D_{1ij} = \sum D_{1ij}^{max} + \sum D_{1ij}^{part} \tag{3.84}$$

$$X_{2j} = \sum^{N} D_{2ij} = \sum D_{2ij}^{max} + \sum D_{2ij}^{part} \tag{3.85}$$

The sums $\sum D_{1ij}^{part}$ and $\sum D_{2ij}^{part}$ are necessarily finite, so as $n \to \infty$, $X_{1j} \to \sum D_{1ij}^{max}$ and $X_{2j} \to \sum D_{2ij}^{max}$ as long as $p_{(DR,\ j)}$, $p_{(DLT,\ j)} > 0$ and the probability of selecting each dose for the next patient is always greater than zero. Thus, $\dfrac{X_{1j}}{n_j} \to p_{(DR,\ j)}$ and $\dfrac{X_{2j}}{n_j} \to p_{(DLT,\ j)}$ by the Weak Law of Large Numbers. As $n \to \infty$, a prediction model is not necessary. If the dose selection mechanism is $max\{x_j \in X :\ \hat{p}(DR,\ j) \leq \pi_{DR},\ \ \hat{p}_{(DLT,\ j)} \leq \pi_{DLT}\}$, then defining $\hat{p}(DR,\ j) = \dfrac{X_{1j}}{n_j}$ and $\hat{p}(DR,\ j) = \dfrac{X_{2j}}{n_j}$ simply, the correct dose x_j^* is guaranteed to be selected.

Chapter 4

Simulations

Simulation is used to compare modeling choices, specifically to identify a best common model, and compare the performance of this common model with that of the TITECRM-MC (Lee et al. 2019), CRM-MC (Lee, Cheng, and Cheung 2011), and TITECRM (Cheung and Chappell 2000). The TITECRM uses only DLT information, while the other two "main" comparator methods also use DLT information as well as Dose Reduction information. Here, only the three-transition illness-death model is studied as it is most relevant. Simulation of the five-transition model is left for future work.

For all simulations, $J = 5$ dose levels are assumed, the first patient starting on the lowest dose. The maximum time of follow-up, as well as the prediction time of interest, is 52 weeks. Sample sizes $n = 30$, $n = 60$, and $n = 100$ are considered, three patients accrued per month. The trial ends when the $(n+1)$ patient is accrued. (With the CRM-MC, this is modified to allow full follow-up for all patients.) The most realistic sample size scenarios are seen as $n = 30$ and $n = 60$, with $n = 100$ included as a large sample type evaluation. For non-constant hazard settings (see below), data is generated via a Cox-Weibull conditional hazard formulation, similar but different from the conditional hazard formulation of the proposed model (for

constant hazard settings, data is generated from the proposed model).

$$h_1 = \alpha\lambda_1 t_1^{\alpha-1} e^{\beta_1 x_j}; \quad S_1 = e^{-\lambda_1 t_1^{\alpha} \exp(\beta_1 x_j)} \tag{4.1}$$

$$h_2 = \alpha\lambda_2 t_2^{\alpha-1} e^{\beta_1 x_j}; \quad S_2 = e^{-\lambda_2 t_2^{\alpha} \exp(\beta_1 x_j)} \tag{4.2}$$

$$h_3 = \alpha\lambda_2 t_2^{\alpha-1} e^{\beta_0+\beta_1 x_j}; \quad S_3 = e^{-\lambda_2 t_2^{\alpha} \exp(\beta_0+\beta_1 x_j)} \tag{4.3}$$

Using the conditional survival functions, data can be generated using inverse transform sampling given $F = 1 - S$, where F is the CDF. Three times to event are generated. If t_1 from S_1 is less than t_2 from S_2, a patient had a Dose Reduction before a DLT; consequently, t_1 and t_2 from S_3 become the two event times, with 52 weeks as a censoring time. If t_2 from S_3 is less than t_1, then no event t_2 is considered possible and 52 is the censoring time. When t_2 from S_2 is less than t_1, the patient has had a DLT before a Dose Reduction, a terminal event, again limited by 52 weeks. This process is valid as the Markov property is assumed, where h_3 and S_3 are measured from time zero.

Data is generated under four hazard rates, or four conditions for α: Constant ($\alpha = 1$), Small Increasing ($\alpha > 1$), Small Decreasing ($\alpha < 1$), and Large Decreasing ($\alpha << 1$). One example of each scenario is shown in Figure 4.1. While the setting assumes chronic treatment and events are expected to occur at late time points, it is expected that relatively more events will occur early in the observation period, making Small Decreasing and Large Decreasing the most likely scenarios in practice.

For each sample size/hazard rate combination, dose-toxicity curves are generated such that the target probability π_{DLT} (0.3 or 0.2) is reached at dose level 1, dose level 3, or dose level 5; and the sister target probability π_{DR} is relatively higher, lower, or equal to its target (0.5 or 0.4) at that same dose. The dose-toxicity curves for $\pi_{DLT} = 0.2$ and $\pi_{DR} = 0.4$ are shown in Figure 4.2. (The curves for $\pi_{DLT} = 0.3$ and

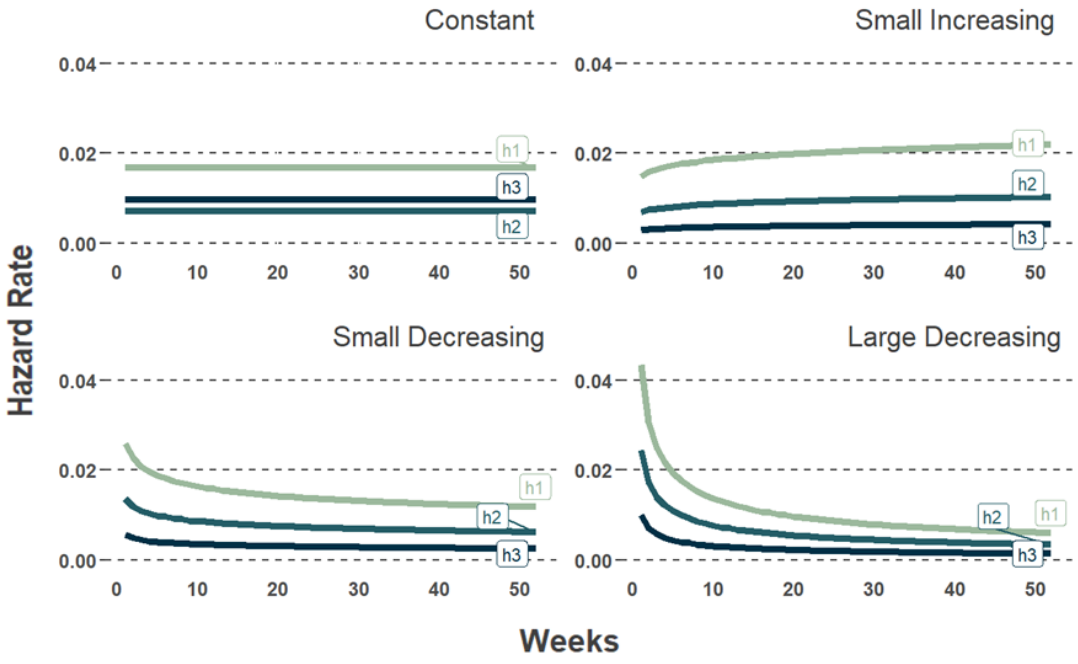

Figure 4.1: Examples of the four hazard rate shapes used to simulate data.

$\pi_{DR} = 0.5$, and parameter settings of all toxicity curves are in Appendix A.)

Models are evaluated based on 1) Percent Correct Selection (PCS), the percentage of time maximum dose that falls under both toxicity thresholds is chosen; 2) Overdose Percentage (OD), the percentage the selected dose is higher than desired; and 3) Accuracy Index for DR and DLT. The Accuracy Index (Cheung 2011) goes beyond PCS, taking into account how "close" the selected dose was to the desired dose,

$$\text{Accuracy index} = 1 - J\ \frac{\sum_{j=1}^{J} \rho_j P(\text{select dose } j)}{\sum_{j=1}^{J} \rho_j},$$

where ρ_j is the absolute difference between the true toxicity probability π_j at dose j and the target probability θ. Just as there are two events (two θs), there are two Accuracy Indices, A_{DR} and A_{DLT}. The range of an Accuracy Index is $[-1, 1]$ where 1

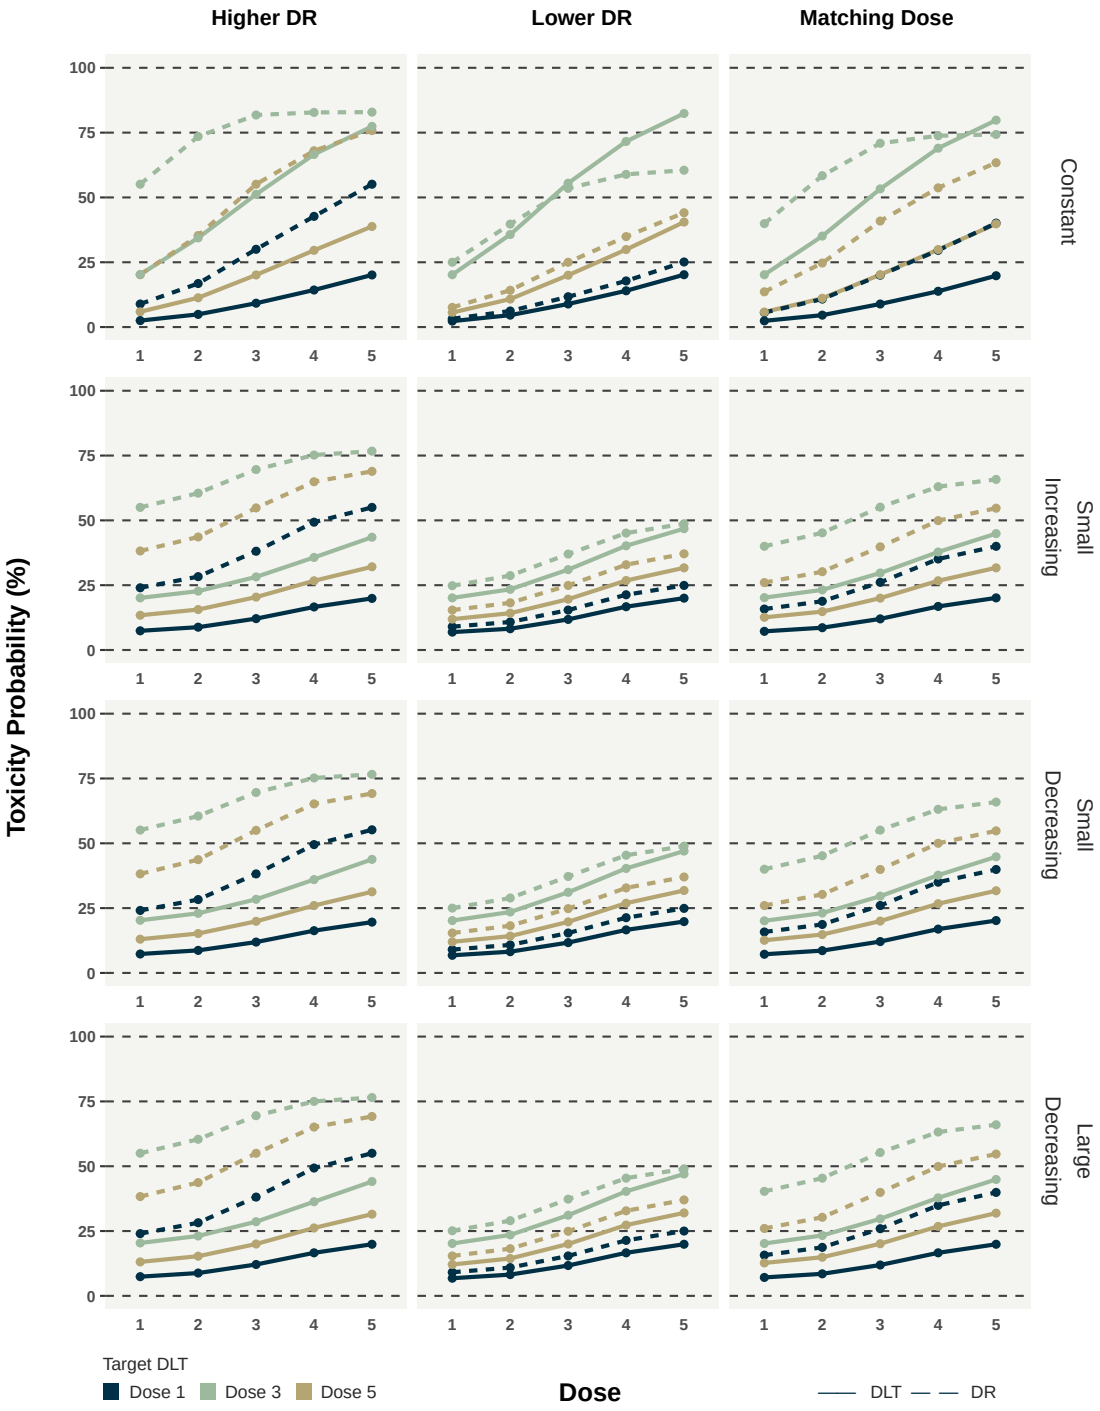

Figure 4.2: Examples of dose-toxicity curves for all data generating hazards ($\pi_{DR} = 0.4$, $\pi_{DLT} = 0.2$).

is perfect selection and zero is equivalent to randomly choosing a dose. All simulation instances are run for $N = 1,000$ iterations. The Bayesian models are fit in STAN within an R environment (Stan Development Team 2022b, Stan Development Team 2022a, R Core Team 2022).

4.1 Model Selection

To select a best common model, the prior distribution, hazard specification, and dosing decision rule are evaluated. One Uniform prior distribution and two Normal prior distributions are considered, one "uninformative" and one "informative," for the constant model only: $(\lambda_1, \lambda_2, \beta_1)$.

Uniform:

$$\lambda_1 \sim \text{Uniform}(0, 0.5) \tag{4.4}$$

$$\lambda_2 \sim \text{Uniform}(0, \lambda_1) \tag{4.5}$$

$$\beta_1 \sim \text{Uniform}(-4, 4) \tag{4.6}$$

Normal (uninformative):

$$\lambda_1 \sim N_+(0.2, 0.2^2) \tag{4.7}$$

$$\lambda_2 \sim N_{+,\lambda_1}(0.1, 0.2^2) \tag{4.8}$$

$$\beta_1 \sim N(0.5, \sqrt{2}^2) \tag{4.9}$$

Normal (informative):

$$\lambda_1 \sim N_+(0.1, 0.05^2) \tag{4.10}$$

$$\lambda_2 \sim N_{+,\lambda_1}(0.05, 0.05^2) \tag{4.11}$$

$$\beta_1 \sim N(0.5, 1^2) \tag{4.12}$$

The "uninformative" Normal prior, in addition to setting higher variances, sets the mean value of baseline hazard rates higher than the "informative" Normal prior, in order to be conservative. A higher baseline hazard rate estimates a higher toxicity at a given dose, consequentially choosing a lower dose. The bounds of the Uniform priors are set just outside realistic values.

Hazard specifications under consideration are the Constant-Skeleton and Weibull-Skeleton as presented in Section 2.1. The method of dosing patients can be different during the trial compared to the end of the trial, thus dosing decision rules are Total Distance, Weighted Total Distance (3:1 DLT to DR), Maximum Uniformly Tolerated Dose, Total Distance during plus MUTD at the end of the trial, and Weighted Total Distance during plus MUTD at the end of the trial, as defined in Section 2.5.

Prior Distribution

Table 4.1 shows dose selection of the five dose levels over nine dose-toxicity scenarios, using the three prior distributions ($n = 30$, with $n = 60$ in parentheses). The Total Distance dosing rule guided dose selection, allowing freedom of movement to better assess the affect of the priors. DR target toxicity and DLT target toxicity ($\pi_{DR} = 0.4$, $\pi_{DLT} = 0.2$ only) occur on the same dose in Scenarios 1, 2, 3, and 6. By the Total Distance criteria, Scenario 4 is split between dose levels 1 and 2, Scenario 5 between dose levels 3 and 4, Scenario 8 between dose levels 2 and 3, and

Scenario 9 between dose levels 3, 4, and 5. Scenario 7 does not describe any doses that satisfy both target toxicities, however for the table values, dose level 1 is considered the desired dose. Any stopping rule for a scenario like this is not considered in the current work.

Dose selection is reasonably similar over priors given each scenario. The informative Normal prior struggles somewhat, especially when target toxicities are on higher doses, most likely because this prior is set where dose is drawn to the middle of the dose set and fights exploration of higher doses. The prior should not restrict exploration of higher doses but should give good probability to all reasonable values. The Uniform prior slightly outperforms the uninformative Normal prior in scenarios where target toxicities fall on the same dose. Where target toxicities differ, the uninformative Normal prior is a bit more conservative. Although it is ultimately the dosing decision rule's responsibility to control dose under its uniformly tolerated bounds, it is desirable to have a prior that allows for exploration of the dose range while not being ignorant of overdosing. Thus, the uninformative Normal prior is selected for use in the rest of the simulation results.

Hazard Formulation and Dosing Decision

Both hazard specification and dosing rule directly affect the flexibility of the method when assigning doses, thus, they are evaluated in tandem. The Weibull-Skeleton model adds a shape parameter α, and another prior is necessary. Fitting with the previous uninformative priors,

$$\alpha \sim N_+(1, 0.3^2). \tag{4.13}$$

Figure 4.3 shows PCS and Overdose percentage for Scenarios 1-3, 5, 8, and 9 for sam-

Scenario		Dose Level					Dose Selection					Prior
		1	2	3	4	5	Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	
1	DR	**40**	58	71	74	74	**92 (97)**	8 (3)	0 (0)	0 (0)	0 (0)	Normal (inf.)
	DLT	**20**	35	53	69	80	**91 (97)**	9 (3)	0 (0)	0 (0)	0 (0)	Normal (un.)
							89 (96)	11 (4)	0 (0)	0 (0)	0 (0)	Uniform
2	DR	14	25	**40**	54	63	8 (0)	60 (29)	**29 (64)**	3 (7)	0 (0)	Normal (inf.)
	DLT	6	11	**20**	30	40	8 (0)	46 (22)	**37 (70)**	8 (8)	1 (0)	Normal (un.)
							8 (1)	42 (19)	**38 (71)**	10 (8)	2 (0)	Uniform
3	DR	6	11	20	30	**40**	0 (0)	15 (0)	43 (5)	28 (39)	**14 (57)**	Normal (inf.)
	DLT	2	5	9	14	**20**	0 (0)	8 (0)	31 (4)	33 (31)	**28 (65)**	Normal (un.)
							0 (0)	5 (0)	29 (3)	37 (30)	**29 (67)**	Uniform
4	DR	25	**40**	54	59	61	**65 (64)**	**34 (36)**	1 (0)	0 (0)	0 (0)	Normal (inf.)
	DLT	**20**	36	56	72	82	**62 (61)**	**35 (39)**	3 (0)	0 (0)	0 (0)	Normal (un.)
							63 (60)	**35 (39)**	2 (1)	0 (0)	0 (0)	Uniform
5	DR	8	14	25	**35**	44	2 (0)	36 (3)	**46 (46)**	**14 (44)**	3 (7)	Normal (inf.)
	DLT	6	11	**20**	30	41	2 (0)	20 (2)	**45 (39)**	**24 (50)**	9 (8)	Normal (un.)
							1 (0)	18 (1)	**48 (39)**	**25 (50)**	8 (9)	Uniform
6	DR	3	6	12	18	**25**	0 (0)	6 (0)	27 (1)	30 (7)	**37 (93)**	Normal (inf.)
	DLT	2	5	9	14	**20**	0 (0)	2 (0)	16 (0)	29 (3)	**54 (97)**	Normal (un.)
							0 (0)	2 (0)	12 (0)	28 (4)	**58 (96)**	Uniform
7	DR	**55**	74	82	83	83	**99 (100)**	1 (0)	0 (0)	0 (0)	0 (0)	Normal (inf.)
	DLT	**20**	34	51	67	77	**99 (100)**	1 (0)	0 (0)	0 (0)	0 (0)	Normal (un.)
							99 (100)	1 (0)	0 (0)	0 (0)	0 (0)	Uniform
8	DR	20	**35**	55	68	76	27 (4)	**60 (75)**	**12 (21)**	1 (0)	0 (0)	Normal (inf.)
	DLT	6	11	**20**	30	39	24 (4)	**56 (70)**	**18 (26)**	2 (0)	0 (0)	Normal (un.)
							24 (3)	**53 (70)**	**20 (26)**	3 (1)	0 (0)	Uniform
9	DR	9	17	**30**	43	55	2 (0)	32 (3)	**46 (35)**	**16 (51)**	**4 (11)**	Normal (inf.)
	DLT	3	5	9	14	**20**	2 (0)	19 (2)	**43 (33)**	**27 (52)**	**10 (14)**	Normal (un.)
							1 (0)	17 (1)	**41 (31)**	**29 (55)**	**12 (13)**	Uniform

Table 4.1: Evaluation of three prior distributions on Constant data generating hazards using Constant model ($n = 30$, with $n = 60$ in parentheses).

ple size $n = 30$. (Remaining scenarios are omitted due to redundancy.) The larger grouping label shown on the right side is the data generating hazard rate. The individual labeling on the left side lists six model/dosing rule combinations: a) Weibull-Skeleton model with 3:1 Weighted Total Distance during the trial + MUTD selection at the end, b) Weibull-Skeleton model with Total Distance during the trial + MUTD selection at the end, c) Weibull-Skeleton model with Total Distance throughout, d) Constant-Skeleton model with Total Distance during the trial + MUTD selection at the end, e) Constant-Skeleton model with 3:1 Weighted Total Distance throughout, and f) Constant-Skeleton model with Total Distance throughout. Note, while Table 4.1 indicates dose levels $(4, 3)$, $(2, 3)$, and $(3, 5)$ are the location of target probabilities in Scenarios 5, 8, and 9, this is a result of constant generating hazards. When the generating hazards are non-constant, the location of target probabilities are instead dose levels $(5, 3)$, $(1, 3)$, and $(3, 5)$, as shown in Figure 4.2.

Scenario 1 is entirely Correct Selection or Overdose, as the desired dose is dose level 1; conversely, Scenario 3 has no overdose as the desired dose is dose level 5. In Scenario 1, the Constant models reveal consistently better PCS, with Constant + MUTD leading the group for each data generating hazard. Both toxicity targets here are met at dose level 1; and the Constant models, especially the Constant + MUTD model, are more conservative. In Scenarios 2 and 3, where the toxicity targets both fall on dose level 3 and dose level 5, respectively, the Weibull models outperform the Constant models. The decreasing generating hazards, considered the most realistic, drastically favor the Weibull models. If data follow a decreasing hazard, but are only observed for a short time period, fitting a Constant model will estimate the average hazard to be higher than it would be over the entire curve. Thus, estimated toxicity is higher, and dose selection is too conservative. The best decision rule for the first three scenarios in terms of PCS is simple Total Distance because both toxicity

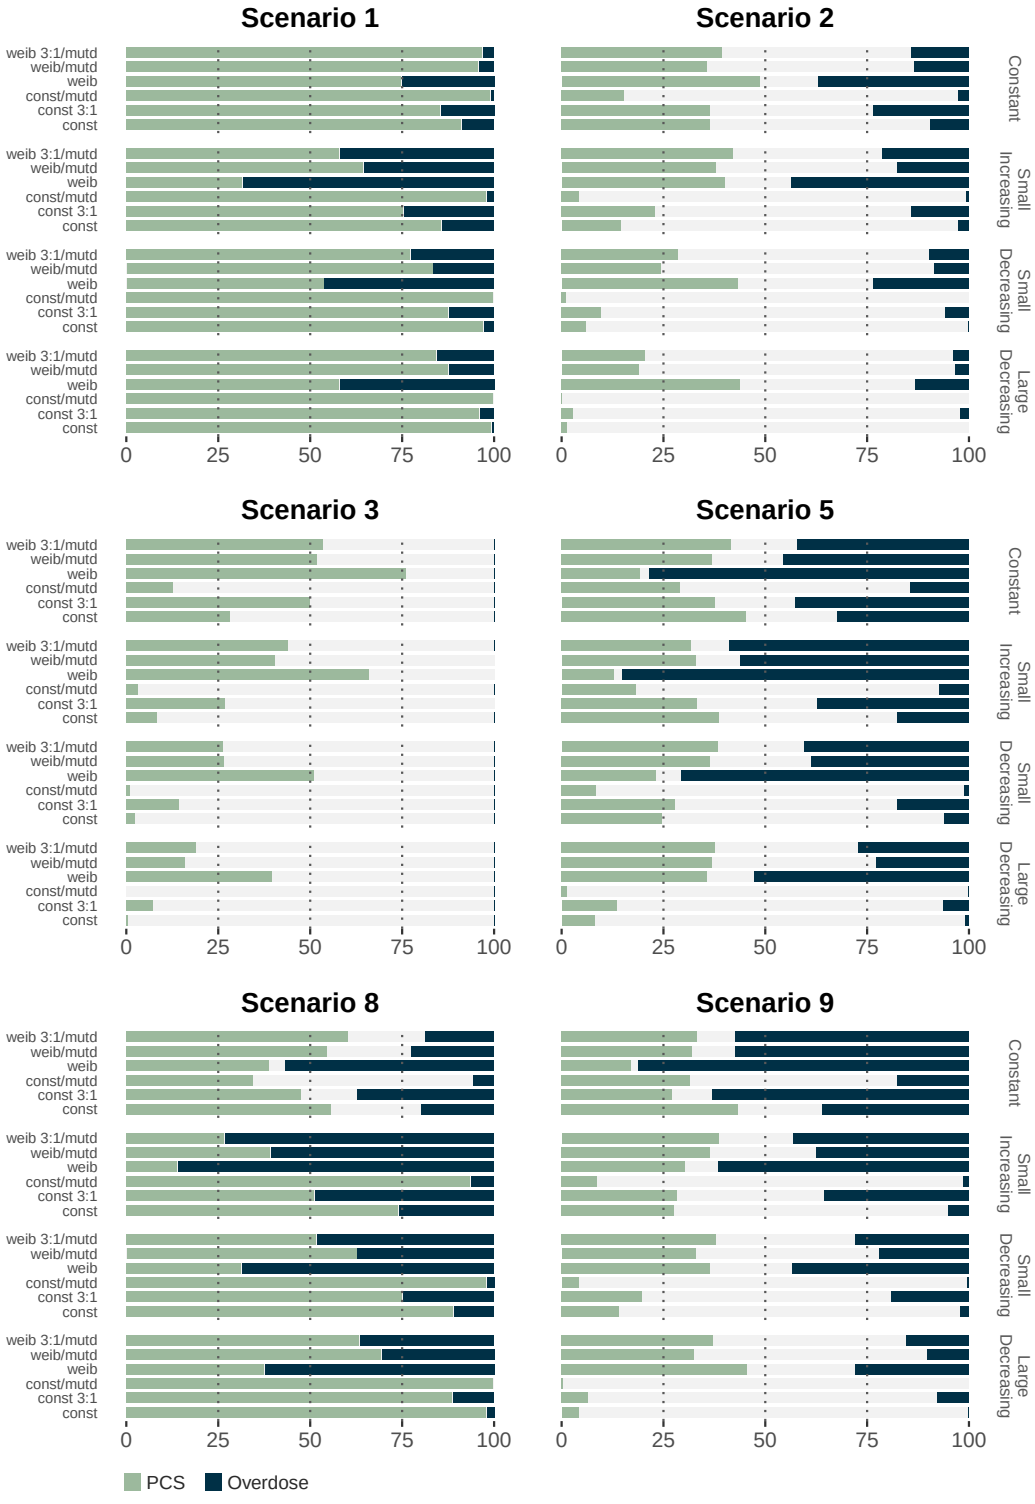

Figure 4.3: Percent Correct Selection (green) and Overdose (navy) for each candidate model combination over all data generating hazards ($n = 30$).

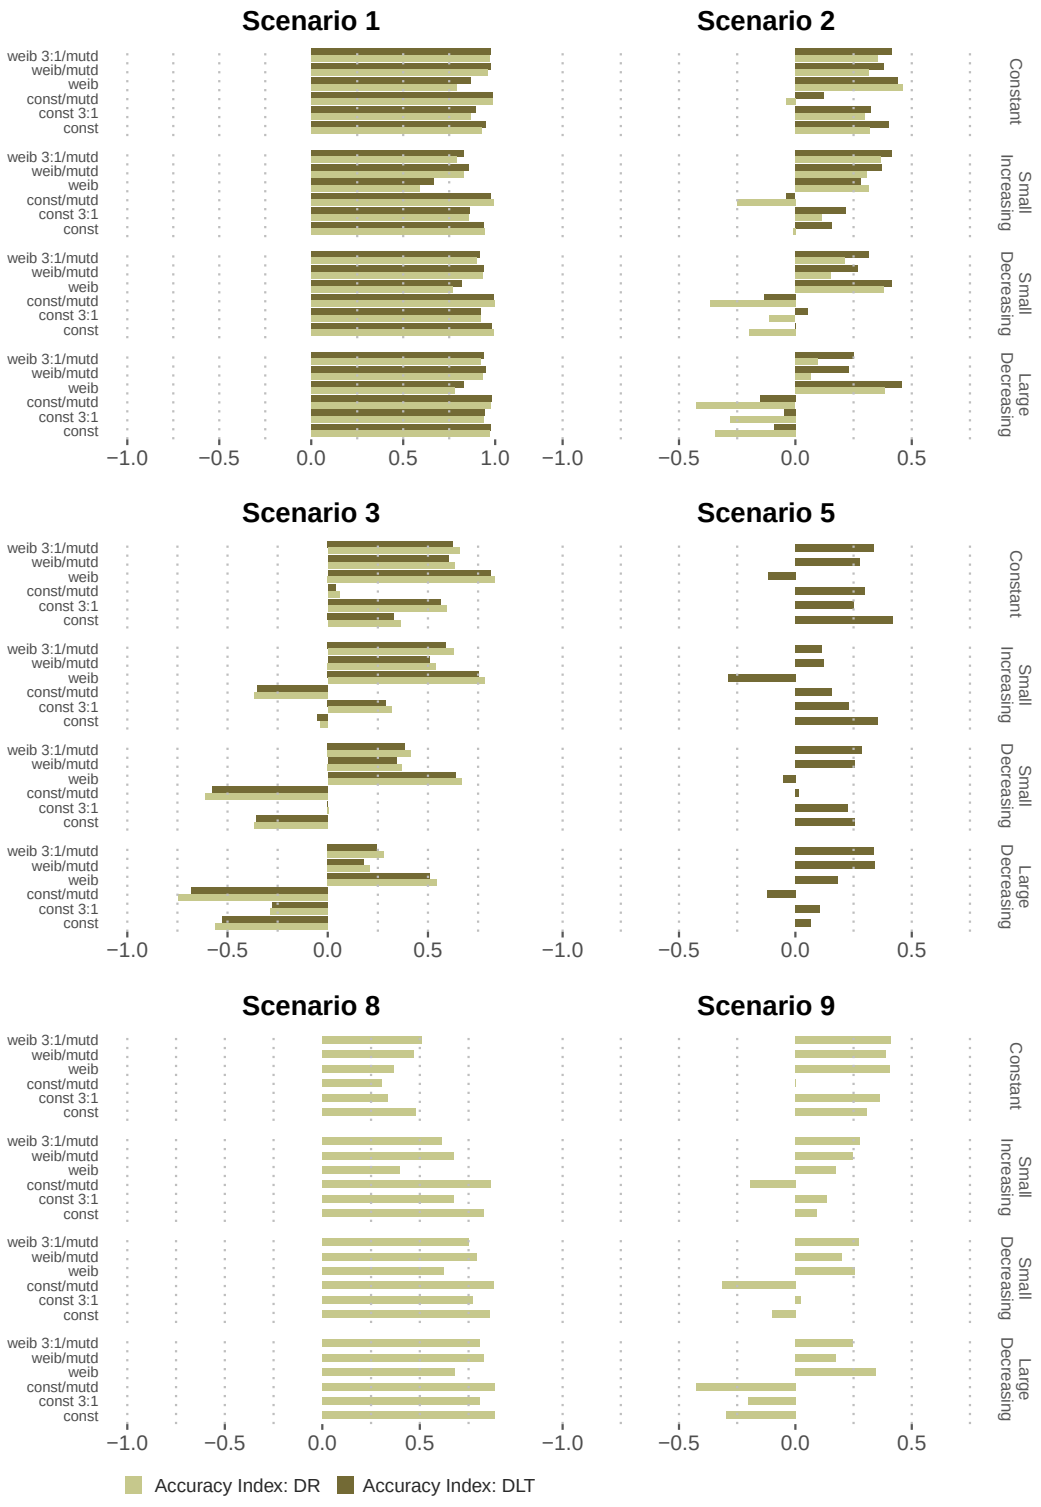

Figure 4.4: Dose Reduction and DLT Accuracy Index values for each candidate model over all data generating hazards ($n = 30$).

targets are leading models to the same dose. However, in Scenario 2 the Weibull + MUTD options fare much better on Overdose. The Total Distance Weibull model sees relatively high Overdose rates.

In Scenario 5, DLT is the limiting toxicity. The current method is less necessary in this potential situation, but because it is a possible scenario and DRs will still occur, it is important to evaluate nonetheless. Under Constant and Small Increasing generating hazards, Weibull + MUTD and Total Distance Constant models perform similarly; however, under decreasing generating hazards, the Weibull models outperform with either Weibull + MUTD option performing best. In Scenarios 8 and 9, DR is the limiting toxicity. These are especially situations of interest which traditional methods would not be able to handle appropriately. Under constant generating hazards, again a Weibull and Constant model are competitive. Under all other hazards, Constant models perform much better than Weibull models, which may seem unexpected; yet the desired dose in all of those scenarios is dose level 1. (DLT toxicity target falls on dose level 3.) This relates back to Scenario 1 where the Constant models are more conservative and will often choose dose level 1. Also, the true probability of DR on dose level 1 ($\sim$38%) and dose level 2 ($\sim$43%) have little separation and are close to the target of 40%. Scenario 9 follows similarly to Scenario 5, with even more drastic underperformance of the constant models. Again, the Weibull + MUTD models present much better Overdose performance than the simple Weibull model.

Figure 4.4 shows the Accuracy Indices for the same scenarios and models. Many of the trends from Figure 4.3 hold here. The Constant models perform better in Scenario 1, but the gap is much smaller. The Weibull models again outperform the Constant models in Scenarios 2 and 3. Only Constant models have instances of negative index, meaning their dose allocation was more inefficient than random dose

selection. Accuracy Index (DR) and Accuracy Index (DLT) mostly mirror each other, except for minor differences in Scenario 2 with decreasing hazards. Scenario 8 has more parity in Accuracy Index than in PCS and Overdose. This provides evidence that the previous explanation of Scenario 8 holds. Constant models are conservative and will stay on dose level 1, while Weibull models stay close to the correct dose level, on average, even if they overdose, choosing dose level 2 instead, just one above. Weibull models again outperform drastically in Scenario 9.

According to these results, it is clear that a Weibull model will fit most scenarios the best and not fall victim to the drastic shortcomings of the Constant models in some scenarios. Additionally, the Weibull + MUTD dosing mechanism performs near the best in most scenarios and has a distinct advantage in overdose control versus the Total Distance Weibull model. Thus, the best common model chosen to move forward is the Weibull model using Total Distance dosing during the trial + MUTD dose selection at the end of the trial, and the uninformative prior specification.

4.2 Model Comparison

Figure 4.5 likewise shows PCS and Overdose Percentage ($n = 30$), here comparing the best common model, Weibull + MUTD, to the three comparator models: TITECRM-MC, CRM-MC (assuming all patients have been followed to full observation), and TITECRM (using partial DLT information only). In Scenario 1, Weibull + MUTD performs similarly to the two main comparators under constant generating hazards, but underperforms for all other generating hazards. This underperformance occurs due to the comparator methods being conservative and choosing dose level 1 more often in any situation. The common model outperforms the comparators in Scenarios 2 and 3 under all generating hazards. In Scenario 5 and 9, the common

model outperforms under non-constant hazards and noticeably more with decreasing hazards. The two main comparator models significantly outperform the common model in Scenario 8. As before, this outcome can be explained, realizing that the comparator models are conservative, and the DR toxicity is tightly spread around the target on dose level 1 as well as dose level 2. The common model has higher overdose percentage in many scenarios even when PCS is higher. Note, TITECRM severely underperforms in Scenarios 5 and 8 as expected, and interestingly Scenario 1 as well.

All models struggle somewhat to reach the highest end of the dose range when dose level 3 or 5 is desired. This is a common problem with dose-finding designs, that is, the desire to prevent overdosing while also exploring the upper range of doses, all with a small sample size. Figure 4.6 shows PCS and Overdose Percentage for $n = 60$. The differences within scenario are similar, with overall more parity between the models. The large sample situation $n = 100$ is left to Appendix A.

Figures 4.7 and 4.8 compare the Accuracy Indices of the best common model and the comparator models. Where the common model was outperformed in Scenario 1 for PCS and Overdose, there is minimal difference in Accuracy Index. The common model remains superior in Scenarios 2, 3, and 9 under all generating hazards. Interestingly, the comparator models are negative or close to zero in most decreasing generating hazard instances for these Scenarios, meaning that on average, dose selection for these models is less efficient than random selection. The common model slightly outperforms the comparators in Scenario 5 under decreasing generating hazards. The drastic difference in Scenario 8 from PCS and Overdose is absent, here realizing only a subtle difference, because the common model stays close to the correct dose even when it overdoses (choosing dose level 2).

Overall, there are positives and negatives to both the Weibull + MUTD and

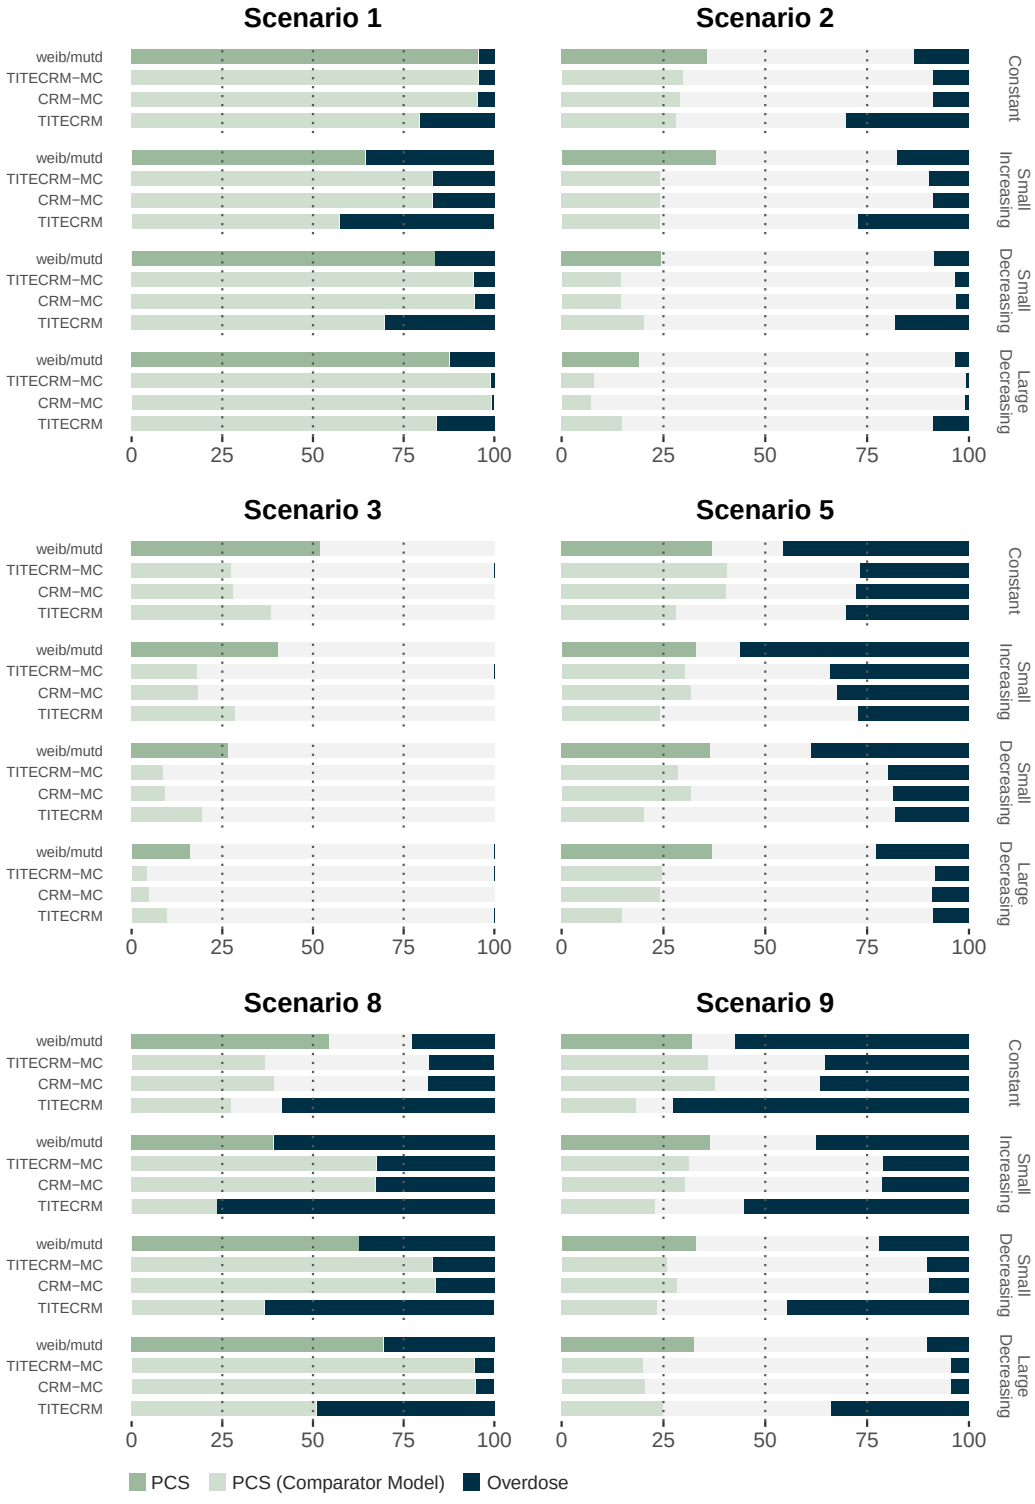

Figure 4.5: Percent Correct Selection (green/light green) and Overdose (navy) for Weibull-Skeleton + MUTD candidate model and comparator models ($n = 30$).

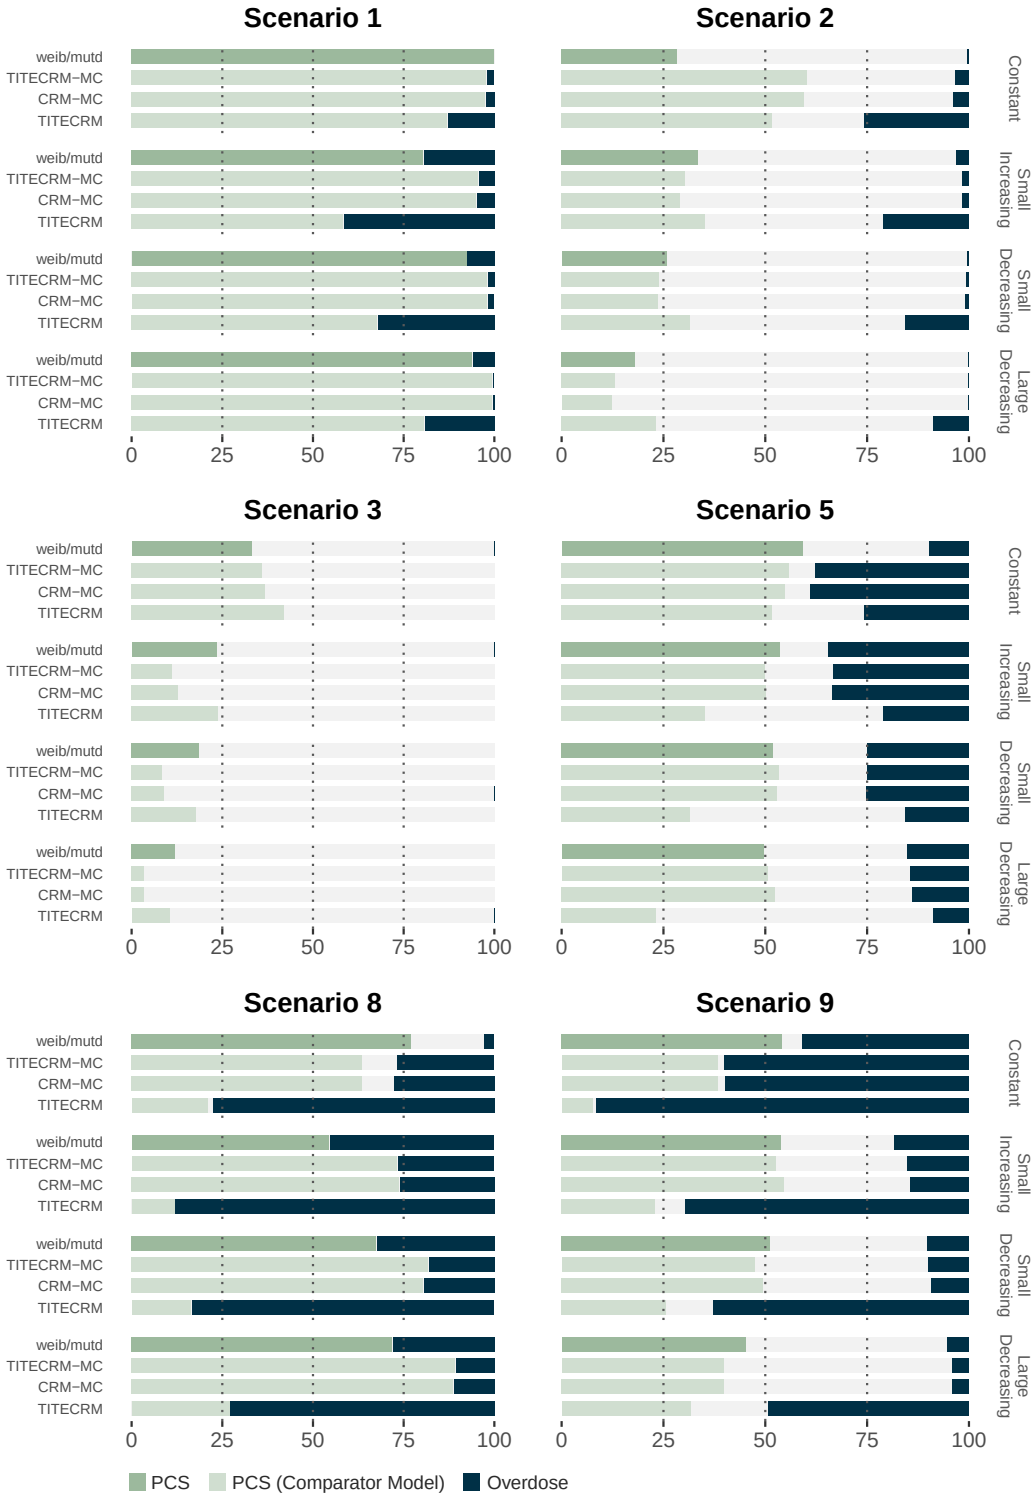

Figure 4.6: Percent Correct Selection (green/light green) and Overdose (navy) for Weibull-Skeleton + MUTD candidate model and comparator models ($n = 60$).

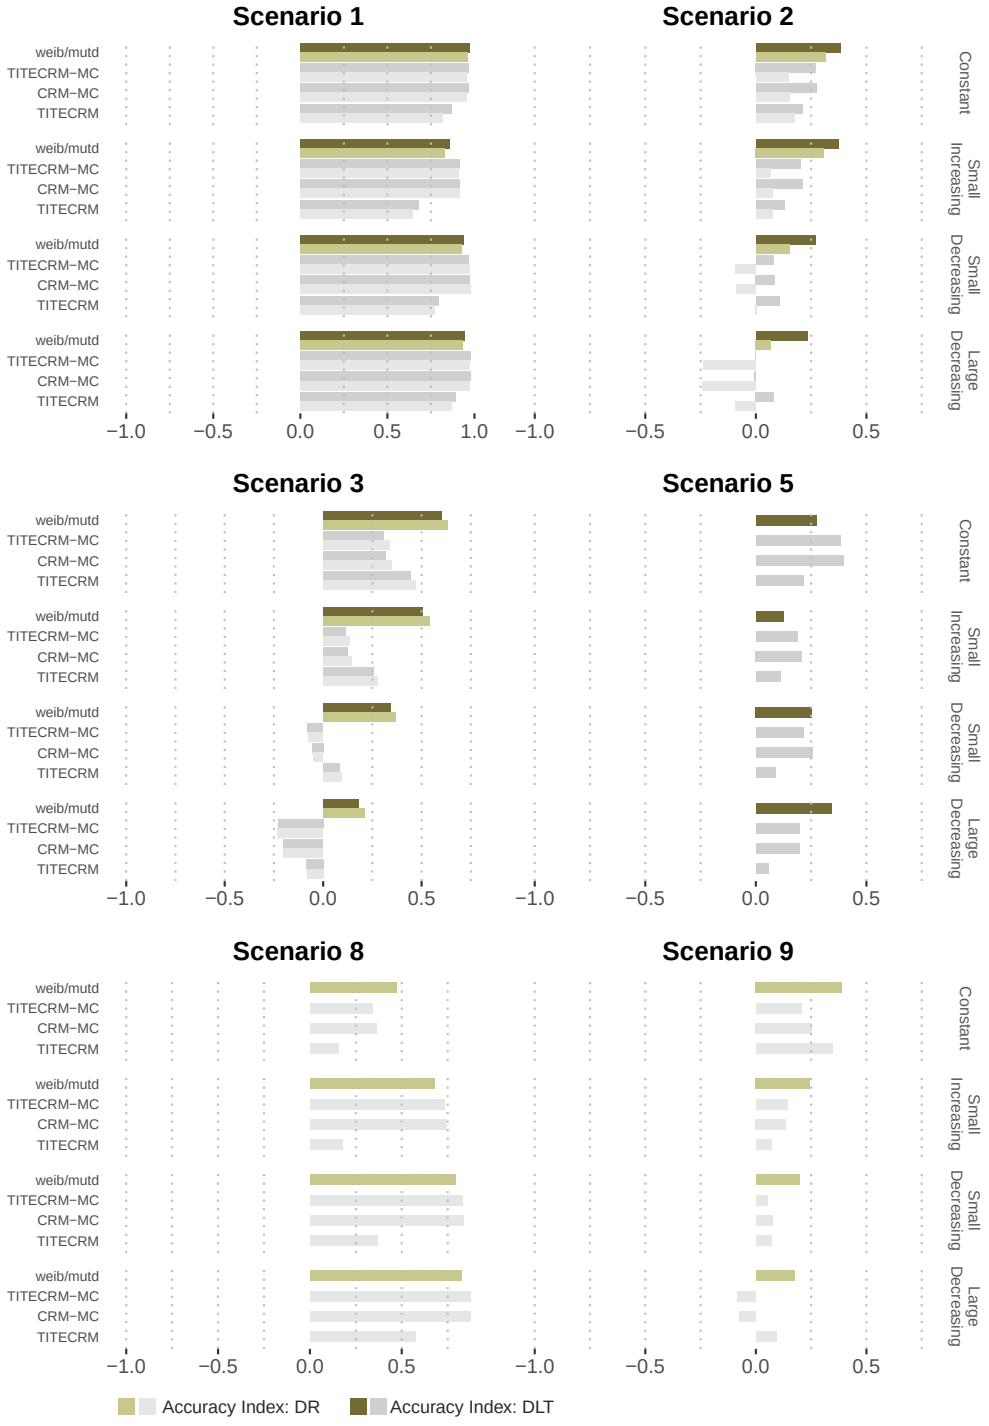

Figure 4.7: Dose Reduction (light brown/light gray) and DLT (brown/gray) Accuracy Index values for Weibull-Skeleton + MUTD candidate model and comparator models ($n = 30$).

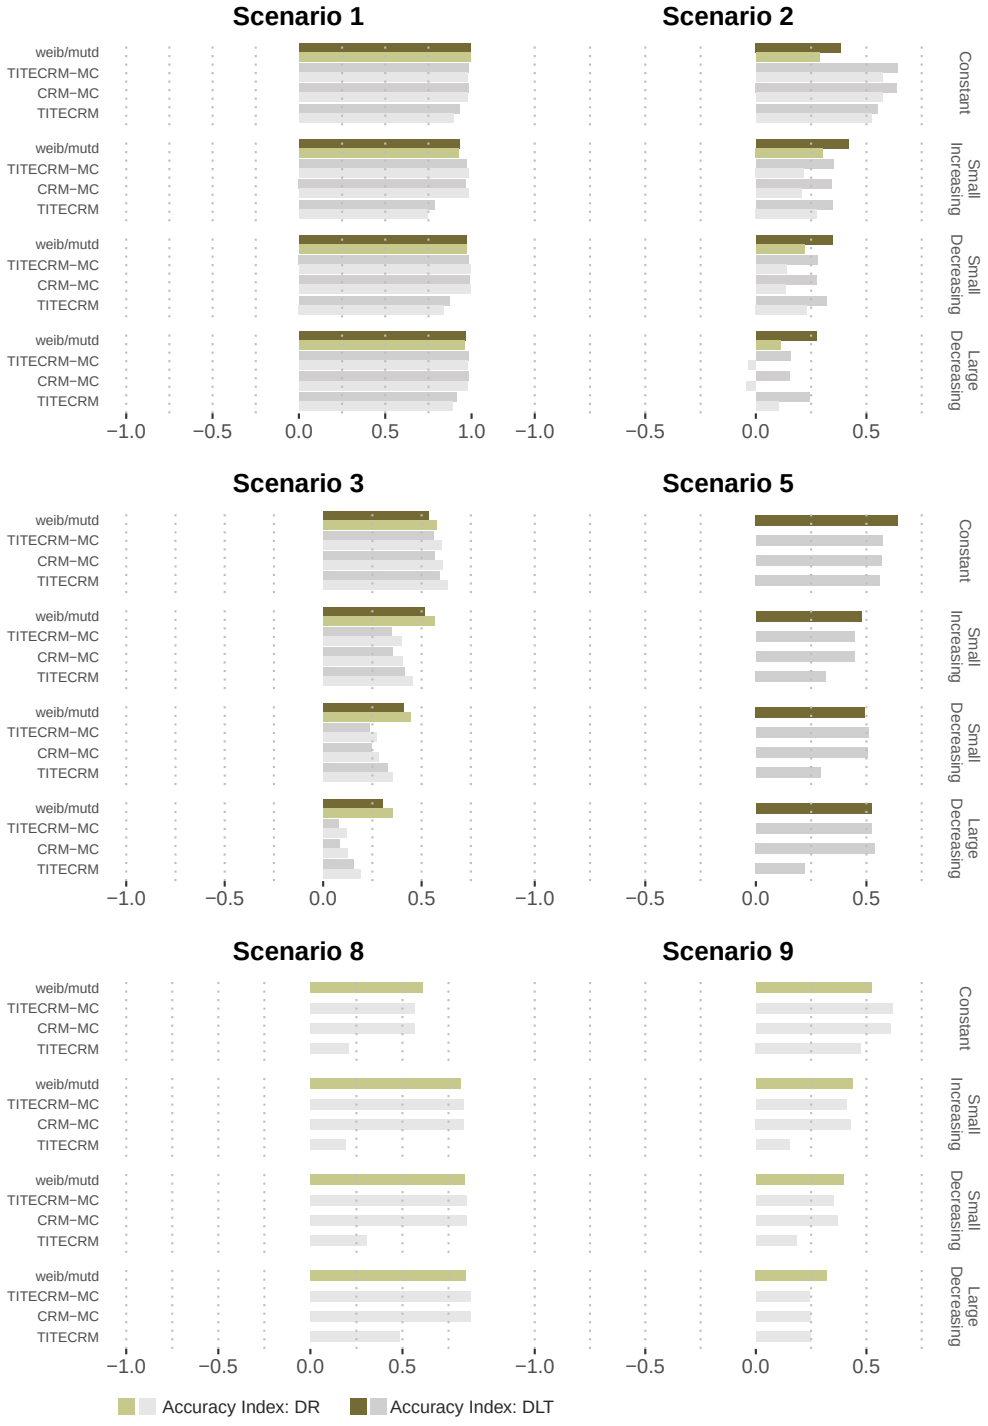

Figure 4.8: Dose Reduction (light brown/light gray) and DLT (brown/gray) Accuracy Index values for Weibull-Skeleton + MUTD candidate model and comparator models ($n = 60$).

TITECRM-MC models. The Weibull + MUTD model excels when the desired dose is higher in the dose range and when the data follow decreasing hazard rates, both of which are most realistic. Additionally, when the Weibull + MUTD model does not select the desired dose, it remains reliably close. There are no instances where the Weibull + MUTD selects dose less efficiently than random, something to which TITECRM-MC falls victim. However, the Weibull + MUTD does overdose more often than the comparators. The most direct remedy for this would likely be a modified decision rule, a path for future work. The TITECRM-MC model excels when desired dose is the lowest in the dose set. It is also fairly competitive with Weibull + MUTD in most instances, and keeps overdose relatively low.

4.3 Finite Sample Size Evaluation

While dose selection is the ultimate arbiter of dose-finding designs, that result is a combination of the underlying model and the dosing decision rule in equal parts. Fundamentally, the proposed toxicity model in this dissertation is designed to estimate probabilities. Bias and variance of these estimates from simulation can help evaluate how well the model is performing at its fundamental task as well as the effect of sample size on this performance. With true probability of toxicity p and estimated probability $\hat{p}$, sample bias, sample variance, and sample mean squared error are defined as

$$bias = (\frac{1}{N}\sum_{i=1}^{N}\hat{p}_i) - p; \tag{4.14}$$

$$variance = \frac{\sum_{i=1}^{N}(\hat{p} - \frac{\sum_{i=1}^{N}\hat{p}}{N})^2}{N-1}; \tag{4.15}$$

$$MSE = bias^2 + variance. \tag{4.16}$$

In the presented simulations, each scenario has a desired dose x_j^*, and two true probabilities of toxicity at that dose, p_{DR} and p_{DLT}. The model generates estimated probabilities $\hat{p}_{DR}$ and $\hat{p}_{DLT}$. Because the proposed model is Bayesian and updates after each patient is accrued, both estimated probabilities are available for any sample size, up to the maximum considered, $n = 100$.

Figure 4.9 shows the bias and the variance of probability estimates on dose x_j^*, and Figure 4.10 shows MSE averaged over the six scenarios shown previously. Again, each of the four data generating hazards, and here both toxicity threshold pairings ($pi_{DR} = 0.4, \pi_{DLT} = 0.2$) or ($pi_{DR} = 0.5, \pi_{DLT} = 0.3$), are included. Both Constant-Skeleton and Weibull-Skeleton models are included. In the proposed method, dosing decision is always based on Total Distance during a trial; so, there is only one instance of each model.

From Figure 4.9, Constant-Skeleton and Weibull-Skeleton models are ultimately unbiased under Constant generating hazards. For all other generating hazards, the Constant-Skeleton model is positive biased over a large sample, while the Weibull-Skeleton model seems to reach the target probability eventually (some rounding of true probabilities likely affects exact unbiasedness), getting quite close by $n = 40$. The Weibull-Skeleton model often underestimates probability at the beginning. The DLT bias is lower for both models compared to DR bias (note the varying y-axis scales on the figure). The variance of probability estimates favors the Weibull-Skeleton model as well. What may seem odd, the Constant-Skeleton model has lower variance at $n = 10$ than $n = 20$ and $n = 30$. However, this outcome is related to the positive bias. The Constant-Skeleton model naturally overestimates hazards and probability, and compounded with a conservative prior and little data, leads to estimates close to 1. In this situation, variance will naturally be low. After the initial jump, variance begins to flatten out a touch later than bias, around $n = 60$.

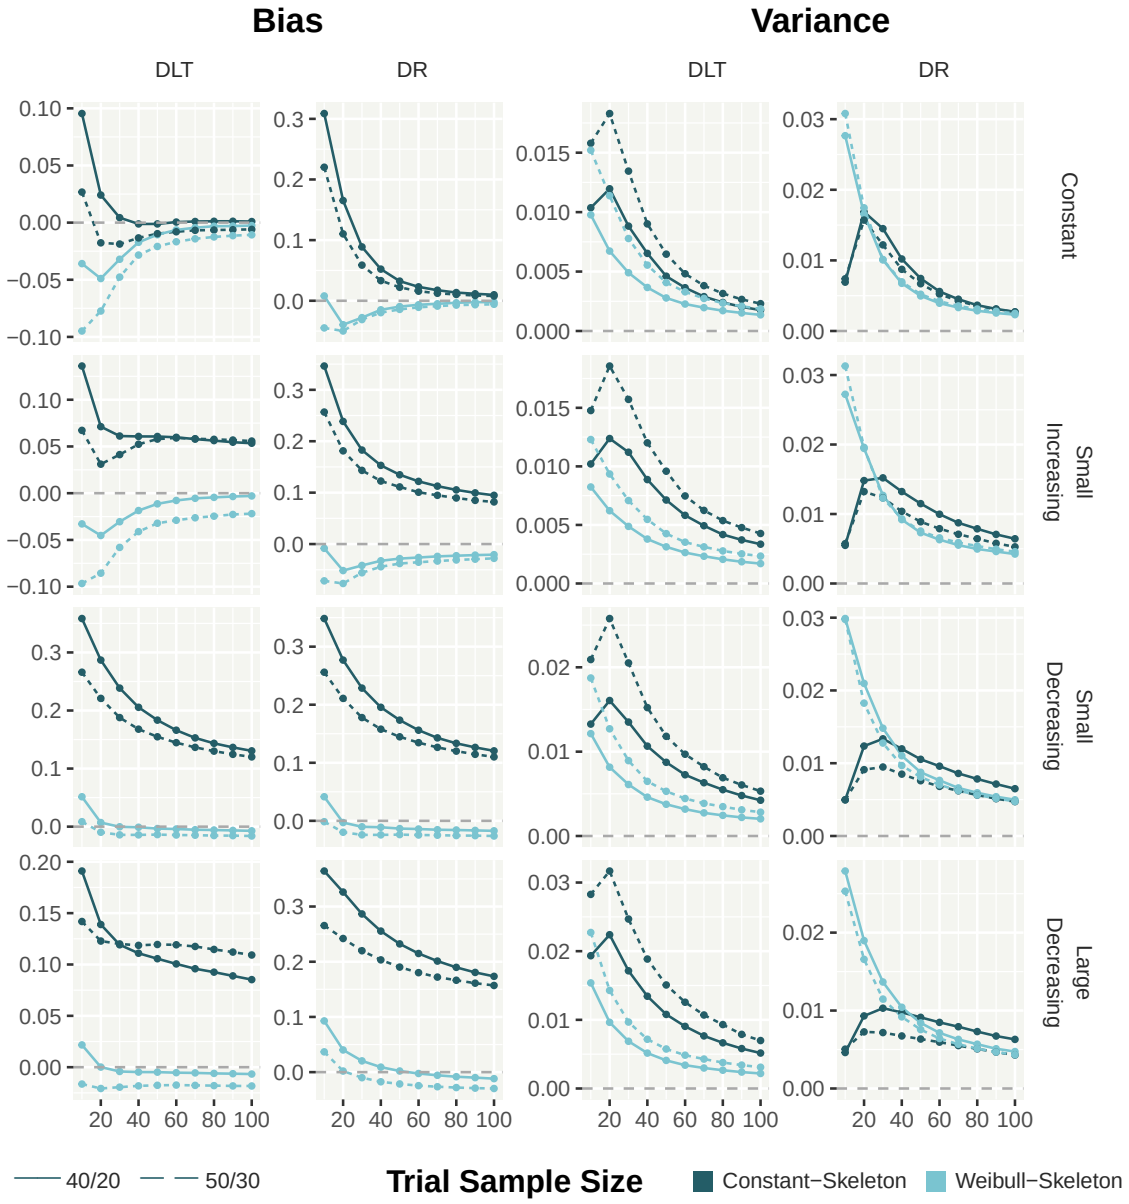

Figure 4.9: Bias and Variance of probability estimates generated by Constant-Skeleton and Weibull-Skeleton models over all data generating hazards and target toxicity thresholds. Average of dose-toxicity Scenarios 1, 2, 3, 5, 8, and 9.

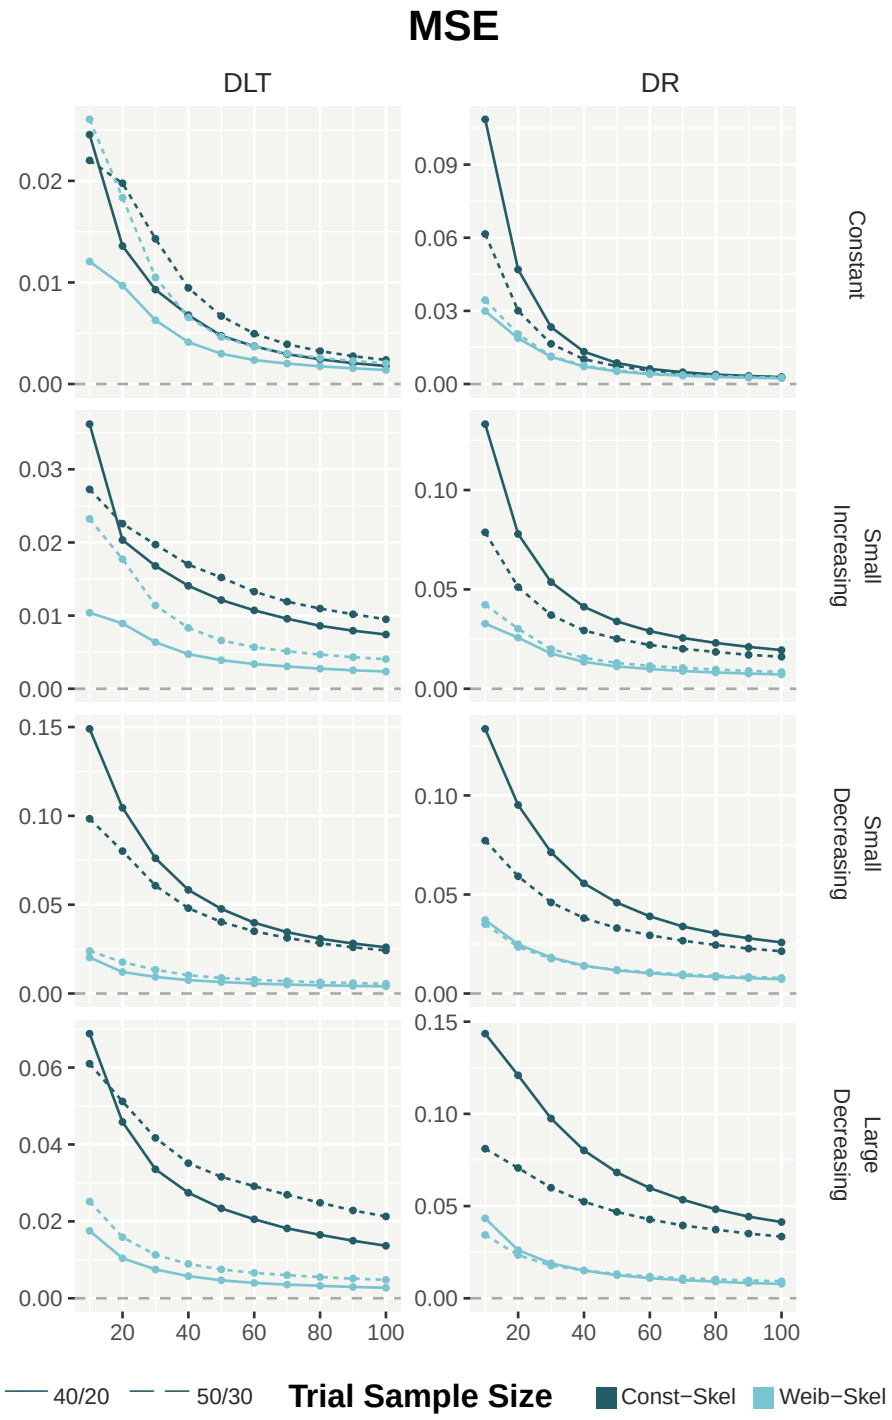

Figure 4.10: MSE of probability estimates generated by Constant-Skeleton and Weibull-Skeleton models over all data generating hazards and target toxicity thresholds. Average of dose-toxicity Scenarios 1, 2, 3, 5, 8, and 9.

The mean squared error, a combination of bias and variance, is monotone decreasing over sample size as expected. The Weibull-Skeleton model is better for DLT and DR probabilities under each generating hazard, even slightly so for Constant generating hazards. Again, $n = 40$ seems to be the sample size where the curve begins to flatten, beyond which extra sample size gives relatively little reduction in MSE.

Overall, bias and variance may be slightly improved in situations when ($\pi_{DR} = 0.4, \pi_{DLT} = 0.2$), but with only subtle differences compared to ($\pi_{DR} = 0.5, \pi_{DLT} = 0.3$). Based on the trajectory of bias, variance, and mean squared error over sample size, and the previous simulation results, a conservative estimate for recommended sample size would be $n = 60$, while $n = 40$ should provide adequate operating characteristics and is recommended for efficiency.

Chapter 5

Efficacy Methods

For Molecularly Targeted Agents and Immunotherapies, efficacy should be evaluated simultaneously along with toxicity in order to recommend doses to carry forward, as efficacy may not be strictly monotone. Yoh et al. 2017, Hamid et al. 2013, Seiwert et al. 2016, and Wolchok et al. 2013 are examples of the common way to evaluate efficacy in an early oncology trial. Tumor burden is recorded at baseline and multiple follow-up time points for each patient. Percent change-from-baseline is calculated, and a responder analysis applied to label Complete Response, Partial Response, Stable Disease, or Progressive Disease. The most common response evaluation criteria is RECIST. Table 5.1 gives RECIST criteria, including definition of tumor burden. Other responder criteria include Modified WHO criteria (WHO) and immune-related Response Criteria (irRC), each quite similar. The percent thresholds in RECIST are required to be confirmed at separate visits ($\geq$ 4 weeks apart), thus it is not strictly a maximum change-from-baseline measure.

The only trial mentioned previously that evaluated more than one dose was Wolchok et al. 2013 in the form of an individual dose escalation trial. The others had already chosen a single dose from a previous trial based on toxicity evaluation. Additionally, in each of these trials, DLTs and Dose Reductions were prevalent; but missing data after DLTs was ignored, and efficacy was associated exclusively with the initial dose the patient received. This is a Treatment Policy/While-on-Treatment strategy for handling DLTs (the strategies are identical if no more follow-up is observed after

Response Category	Definition
Complete Response (CR)	Disappearance of all target lesions.
Partial Response (PR)	At least 30% decrease in the sum of the Longest Diameter (LD) of target lesions, taking as reference the baseline sum LD.
Stable Disease (SD)	Neither sufficient shrinkage for PR nor increase for PD.
Progressive Disease (PD)	At least 20% increase in the sum of the LD of target lesions, taking as reference the smallest sum LD recorded since the treatment started or the appearance of one or more new lesions.

Table 5.1: RECIST responder criteria

DLT) and a Treatment Policy strategy for Dose Reductions, where DLTs and DRs are intercurrent events in the estimand framework (see Chapter 1 for definitions). Treatment Policy is the primary estimand strategy for efficacy trials in the United States as it is meant to provide a realistic description of efficacy, including events that are expected to occur in practice. The estimand resulting from a combination of these strategies and a maximum change-from-baseline-type outcome, defined fully in Table 5.2, attempts to answer the question, What is the maximum achievable tumor response for a given dose in a clinical setting?

Element	Characteristic
Population	Patients with certain type of cancer, inclusion criteria.
Variable	Maximum CfB or Maximum CfB (2 visits) by time t (e.g. 48/52 weeks).
Intercurrent Events	Dose Reduction (Treatment Policy), DLT (Treatment Policy/While-on-Treatment).
Population Level Summary	Proportion of overall responders, by dose.

Table 5.2: Estimand definition of common maximum change-from-baseline outcome criteria.

However, chronic therapies, such as targeted agents and immunotherapies, should be evaluated chronically, characterizing the full efficacy profile of the drug. This

estimand, then, fails to identify the treatment effect of interest. A more relevant question would be, What is the expected efficacy by time t after an expected number of toxicity events? The next section presents an illustration of artificial data to help describe these shortcomings and present a better estimand.

5.1 Illustration

Tumor burden data from the previously mentioned trials (Figure 1.2 and Appendix A.1) are valuable for observing the possible trajectories and levels of noise in percent change-from-baseline response. However, the same data are limited in their ability to study potential estimands as they do not label which patients have had DLTs, label which patients have had Dose Reductions and when, or contain multiple initial dose groups. Thus, an illustration of artificial data is generated in Figure 5.1 where all of this information is known.

Three dose levels are assumed, with six patients on each dose measured at baseline as well as every 8 weeks up to 48 weeks. Patients can come off the trial early for an uninformative reason (large circle) or a DLT (large square), and can have a Dose Reduction while continuing the trial (large triangle). A dose level 0 is assumed to exist to which dose level 1 patients would be reduced. On dose level 3, 2/6 patients have a DLT and 3/6 patients have a DR, with 1/6 DLTs and 2/6 DRs on dose level 2, and 0/6 DLTs and 1/6 DRs on dose level 1. All dose levels have two patients with missing data for reasons unrelated to either the toxicity or efficacy outcomes of the trial, such as loss to follow-up or trial end.

The “Small Regression,” “Medium Regression,” and “Large Regression” scenarios refer to the degree to which patients tend to regress (increase) after a period of tumor reduction, as evidenced in practice in Figures 1.2 and Figure A.3. The responder

cutoffs are set at 80%, 30%, and 0% reduction (a slight modification of RECIST). In this situation, assume the drug development stakeholders are looking for at least a 50% overall response rate (PR + CR) to move forward with a dose. Toxicity thresholds are set as $\pi_{DR} = 0.5$ and $\pi_{DLT} = 0.33$.

The three scenarios have noticeably different average trajectories. The Small Regression scenario exhibits monotonic tumor change almost exclusively, with most reduction occurring early, similar to Figure A.2. Conversely, most patients in the Medium Regression scenario have a non-monotone U-shaped trajectory, where tumor burden increases by the end of follow-up after an initial sharp decrease, similar to Figure A.3. The Large Regression scenario is a more extreme U-shaped example, similar to a version of Figure 1.2, with longer observation. If the entire trajectory is important in assessing efficacy, as asserted, intuitively these scenarios have differing efficacy profiles over a follow-up time of interest 48/52 weeks. However, Table 5.3 shows that the traditional estimands would describe these scenarios as exactly or almost exactly the same.

The Max CfB estimand has the same number of responders (PR and CR) in each scenario for a given dose. The amount of regression in later follow-up visits makes no difference. The Max CfB (2) estimand, meant to represent a RECIST-type evaluation, is almost exactly the same over scenarios, with a few minor differences. In the Medium and Large Regression scenarios, Max CfB (2) assigns one overall responder on Dose 1 (instead of two) and four overall responders on Dose 3 (instead of five). Using just sample estimates, all three doses are safe (Dose 3 just on the threshold), and Doses 2 and 3 are efficacious on either estimand. These two doses could move forward.

The traditional estimands are not able to use the observed level of regression to meaningfully distinguish between scenarios. Similarly, the effect of dose reductions

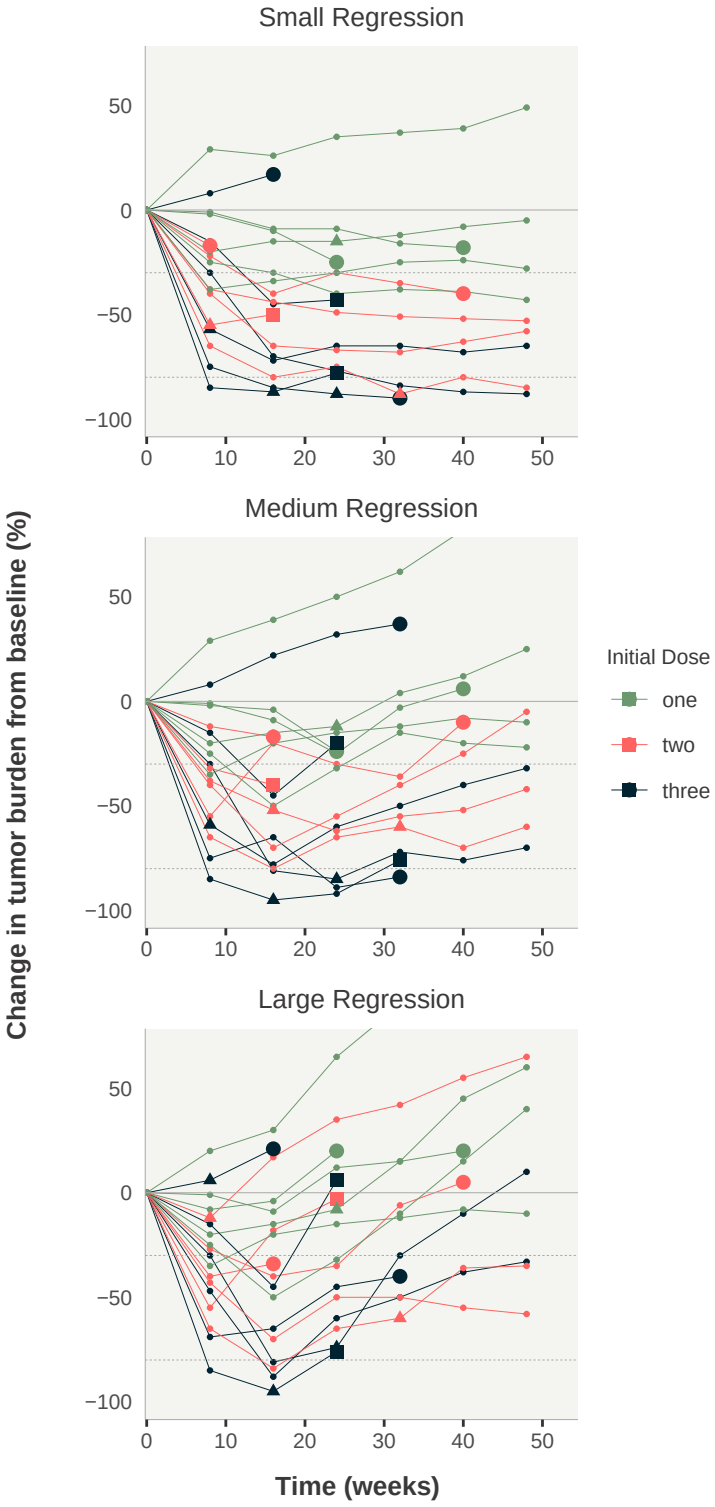

Figure 5.1: Artificial percent tumor change data for three different trajectory situations including missing data.

Dose	Estimand	Scenario														
		Small Regression					Medium Regression					Large Regression				
		Efficacy					Efficacy					Efficacy				
		prog.	stable	part.	comp.		prog.	stable	part.	comp.		prog.	stable	part.	comp.	
Dose 1	Max. CfB	1	3	2	0	**2**	1	3	2	0	**2**	1	3	2	0	**2**
	Max. CfB (2)	1	3	2	0	**2**	1	4	1	0	**1**	1	4	1	0	**1**
	LOCF	1	4	1	0	**1**	3	3	0	0	**0**	5	1	0	0	**0**
Dose 2	Max. CfB	0	1	4	1	**5**	0	1	4	1	**5**	0	1	4	1	**5**
	Max. CfB (2)	0	1	4	1	**5**	0	1	5	0	**5**	1	1	4	0	**4**
	LOCF	0	1	4	1	**5**	0	3	3	0	**3**	3	0	3	0	**3**
Dose 3	Max. CfB	1	0	2	3	**5**	1	0	2	3	**5**	1	0	2	3	**5**
	Max. CfB (2)	1	0	2	3	**5**	1	1	1	3	**4**	1	1	2	2	**4**
	LOCF	1	0	3	2	**5**	1	1	3	1	**4**	2	1	3	0	**3**

Table 5.3: Responder analysis of artificial data using traditional estimands.

is not taken into account. If a patient experiences a dose reduction, their efficacy trajectory will almost certainly be diminished compared to the patient remaining on their initial dose. But because the maximum tumor decrease often occurs in one of the first visits, the effect of a dose reduction will not be reflected much in this measure, if at all. The effect would show in later visits.

A missing data method common in efficacy trials uses the last observation for each patient, or Last Observation Carried Forward (LOCF), again recording post-DLT visits as missing and associating the response with initial dose. An estimand using this method is defined in Table 5.4. The Population and Intercurrent Event elements remain the same, but the Variable and Population Level Summary elements are now redefined. LOCF is the most basic way to estimate the response at time of interest t in the face of missing data. The efficacy responses using this estimand are shown in Table 5.6 as well (LOCF).

Element	Characteristic
Population	Patients with certain type of cancer, inclusion criteria.
Variable	Change-from-baseline at time t (e.g. 48 weeks).
Intercurrent Events	Dose Reduction (Treatment Policy), DLT (Treatment Policy/While-on-Treatment).
Population Level Summary	Proportion of overall responders, by dose.

Table 5.4: Estimand definition using Last Observation Carried Forward method.

The LOCF estimand attempts to answer a more appropriate question, What is the efficacy by time t for a given dose in a clinical setting? It can incorporate some tumor regression and some effect of dose reduction, characterizes the observed trajectory, and would be preferred over the previous estimands. However, the LOCF estimand fails to answer its question appropriately. It suffers from a time scale issue, again caused by the potential U-shaped trajectory. If a dose level contains a high number of DLTs happening near the bottom of the "U," these DLT patients will be recorded as experiencing a large reduction in tumor burden because the observation is occurring on an abbreviated time scale. The LOCF procedure implies the response profile of the subject remains constant after missing until the timepoint of interest. A lower dose that causes minimal DLTs would be seen as much worse because these patients complete follow-up and follow their U-shape through its regression, even if maximum tumor change between doses are similar. The more toxic dose benefits from causing a harmful event. This would not be an issue if efficacy trajectory were reliably monotonic. In this case, coming off the drug early would express poorly in the recorded efficacy, and would be a realistic representation of what to expect in a clinical setting, so the LOCF estimand would suffice. Significant literature exists criticizing the shortcomings of LOCF endpoints (Gadbury, Coffey, and Allison 2003, Streiner 2008, Lachin 2016).

Increasingly more often, an estimand procedure is being used in clinical trials where negative intercurrent events are incorporated into the efficacy response variable itself via "poor" imputation, in an attempt to address multiple outcomes with a single measure. In this procedure, patients who experience an intercurrent event and subsequently produce missing data are not imputed with their expected outcome, but with an artificial poor outcome worse than expected. This is a Composite estimand. With the proposed model, Dose Reductions are included in the toxicity evaluation, thus a Composite estimand is unnecessary. What is desired is a realistic outlook for efficacy, while addressing the issues at hand.

5.2 Two-Level Multiple Imputation

With the potential of U-shaped tumor burden trajectories, data from the latter half of follow-up must be considered in order to fully characterize the efficacy profile of the dose/drug. Moreover, a common time point is necessary to properly compare efficacy among doses. Patients with incomplete data should have their missing visits imputed by a model, one that incorporates a potential U-shape, (doses can then be compared at a common time point), as well as Dose Reductions and DLTs (to include the effect of toxicities). The goal is to create an estimand that follows the theme of Treatment Policy while addressing the major issues caused by missingness.

Multiple imputation (M.I.) is a common and accepted procedure for handling missing data. Covariates can be included that better inform imputation of the response, validate a Missing at Random assumption, and can be imputed themselves if missing. The present setting is rare in the M.I. context, though, because longitudinal data, like tumor burden, is two-level or clustered data. Multiple responses are recorded on the same patient over time, thus each response is not independent;

responses within an individual are correlated. A two-level multiple imputation model is necessary. Initial dose, DLT occurrence, and Dose Reduction occurrence are important covariates in the current setting. Other covariates could be included as well, if measured. The estimand definition for this Treatment Policy M.I. estimand is given in Table 5.5.

Element	Characteristic
Population	Patients with certain type of cancer, inclusion criteria.
Variable	Change-from-baseline at time t (e.g. 48 weeks).
Intercurrent Events	Dose Reduction (Treatment Policy with M.I.), DLT (Treatment Policy with M.I.).
Population Level Summary	Proportion of overall responders, by dose.

Table 5.5: Estimand definition for the proposed Multiple Imputation-based method.

Assumptions

By definition, data is monotone missing for a patient after a DLT. An imputation model thus makes an assumption about patients after they have a DLT. For example, if initial dose is included as the only covariate in a two-level imputation model, this assumes DLTs and Dose Reductions are uninformative. Patients with DLTs and Dose Reductions are expected to have the same average trajectory as those with no toxicity events, which is unrealistic. The goal is to make an assumption that most closely follows the theme of Treatment Policy, or what would be expected to occur in practice. What happens after a DLT in practice can vary. A patient may likely begin a different therapy, or may come off therapy entirely. It is also possible that DLTs could be reversible, and a patient could return back to the drug after a period of time, akin to a more severe Dose Interruption or Reduction.

If DLTs are reversible, then a patient with a DLT could be assumed to move to the next lower dose level and imputed. Toxicity occurrence could be imputed as well to see if any more dose reductions should be expected during the missing visits. If a patient is assumed to come off therapy entirely, this assumption could be implemented directly as a "jump-to-reference," where the patient is imputed based on a reference or placebo group. To assume a patient goes on a different therapy is undesirable, even if expected, because information regarding the different therapy for use in the imputation model might be hard to come by, and the resulting efficacy estimates would be a combination of effects between the two therapies. In this case, it is preferable to again assume a "jump-to-reference" group. The effect of DLTs would be appropriately incorporated, and efficacy attributable to the drug of interest would be isolated.

Information on a placebo group is preferred in order to implement jump-to-reference in the imputation model, signaling a switch to placebo when a DLT occurs. However, most dose-finding trials do not include a placebo group. A substitute could be imputing DLT patients based on switching to the lowest dose given in the trial. Theoretically, this action should overestimate efficacy, but would still be closer to the true effect of interest than traditional estimands.

For the current illustration, no placebo data is assumed to be available, so jump-to-reference via dose level 1 is the best option. Any patient who has a DLT is assumed to move to dose level 1 for the remainder of their visits. Patients who come off for an uninformative reason have both dose and tumor burden imputed. This imputation model then begs the question, What is the estimated efficacy of patients at time t, attributable to the drug of interest, after an expected number of toxicity events? Dose Reductions are also treated with Treatment Policy by including their occurrence in the imputation model. Dose Reductions are an event expected to occur in practice,

and the question "what is a patient's expected efficacy if they stayed on their initial dose?" is not of interest.

Model

Van Buuren and Groothuis-Oudshoorn 2011 implement a Bayesian Gibbs sampler for the multivariate linear mixed effects model (Schafer and Yucel 2002) in the special case where within-group variance is homogeneous. Let the multivariate linear mixed effects model be represented as

$$y_j = X_j\beta + Z_j b_j + \epsilon_j, \tag{5.1}$$

where y_j is an $n_j \times r$ matrix of multivariate responses from j clusters (individuals in this case), $j = 1, ..., m$. The matrices $X_j(n_j \times p)$ and $Z_j(n_j \times q)$ are known covariates, and $\beta(p \times r)$ and $b_j(q \times r)$ are fixed effects and random effects, respectively. The n_j individual response vectors are independently distributed $N(0, \Sigma)$ while each vector of b_j is $N(0, \Psi)$.

Assume parts of $y_j = (y_1, y_2, ..., y_m)$ are Missing at Random. Let $y_{j(obs)}$ and $y_{j(mis)}$ denote the observed and missing parts of y_j, respectively, and let the unknown parameters be $\theta = (\beta, \Sigma, \Psi)$. Multiple Imputation generates m independent draws of $y_{j(mis)}$ from the posterior predictive distribution for the missing data,

$$P(Y_{j(mis)}|Y_{j(obs)}) = \int P(Y_{j(mis)}|Y_{j(obs)}, \theta)P(\theta|Y_{j(obs)})d\theta, \tag{5.2}$$

where $P(\theta|Y_{j(obs)})$ is the posterior distribution. See Van Buuren and Groothuis-Oudshoorn 2011 for more details regarding the Gibbs sampling and imputation algorithm to estimate $P(Y_{j(mis)}|Y_{j(obs)})$ and sample $y_{j(mis)}$.

The response variable tumor burden is imputed as a continuous variable with time, time^2, and dose level as fixed effects and time, time^2, and an intercept with additional random effects. Including dose level itself to signify toxicity events means that Dose Interruptions, if recorded, would not show up in the model (regular Treatment Policy strategy). This assumption is reasonable as Dose Interruptions are not expected to meaningfully affect the efficacy trajectory as the patient stays on their initial dose.

Depending on the assumption made of patients post-DLT, dose level may be missing as a covariate as well. Dose level is a categorical variable with more than two categories. Without a two-level multinomial model available, a single-level multinomial model is used to impute missing dose level values with time, time^2, tumor burden, and initial dose level as covariates. The illustrated data is imputed via the MICE package in R, using a Fully Conditional Specification (FCS) approach, where missing covariates and missing responses are imputed through specifying a conditional density for each incomplete variable, via the models described previously. The imputation algorithm can be summarized as follows.

1. Create random starting imputations for missing tumor response and dose level covariate.

2. Fit dose level model with its "complete" covariates. Generate imputations for missing dose level.

3. Fit tumor burden model with its "complete" covariates, using previously imputed values of missing dose level. Generate imputations for missing tumor burden.

4. Fit analysis model if desired.

5. Repeat steps 2 through 4 m times.

If desired, any valid full-data analysis model could be fit to the imputed data at each step, and the resulting parameter estimates combined via Rubin's Rules (Rubin 2004). For the purpose of this illustration, imputed data alone is enough. Average imputed tumor burden data from the two-level multiple imputation model with $m = 20$ imputations is shown in Figure 5.2.

Illustration results

Patients who initially experienced tumor reduction are imputed with a seemingly reasonable level of regression (increase) based on the scenario. Those patients with an initial increase in tumor burden show a regression to the mean that may or may not be realistic in all scenarios, but on average should still allow for effective conclusions.

Table 5.6 shows the outcomes of previous estimands with the Imputation estimand now included. Dose 2 and Dose 3 Imputation outcomes are slightly different over scenario compared to the LOCF estimand. The important changes result from Dose 2 in Medium Regression and Dose 2 and 3 in Large Regression. In the Medium Regression scenario, all estimands describe Dose 3 as efficacious and Dose 1 as non-efficacious using sample estimates. With Dose 2, the three previous estimands all show efficacy (5/6, 5/6, 3/6 overall responders) while Imputation does not show efficacy (2/6). According to the three previous estimands, Dose 2 and Dose 3 could be selected to move forward, but the Imputation estimand would choose only Dose 3.

In the Large Regression scenario, the differences between estimands exacerbate. Dose 2 offers the same conclusion as in the Medium Regression scenario, with Imputation as the only estimand showing no efficacy (2/6). Here, Imputation shows Dose 3 as non-efficacious as well (2/6), while each of the previous estimands shows

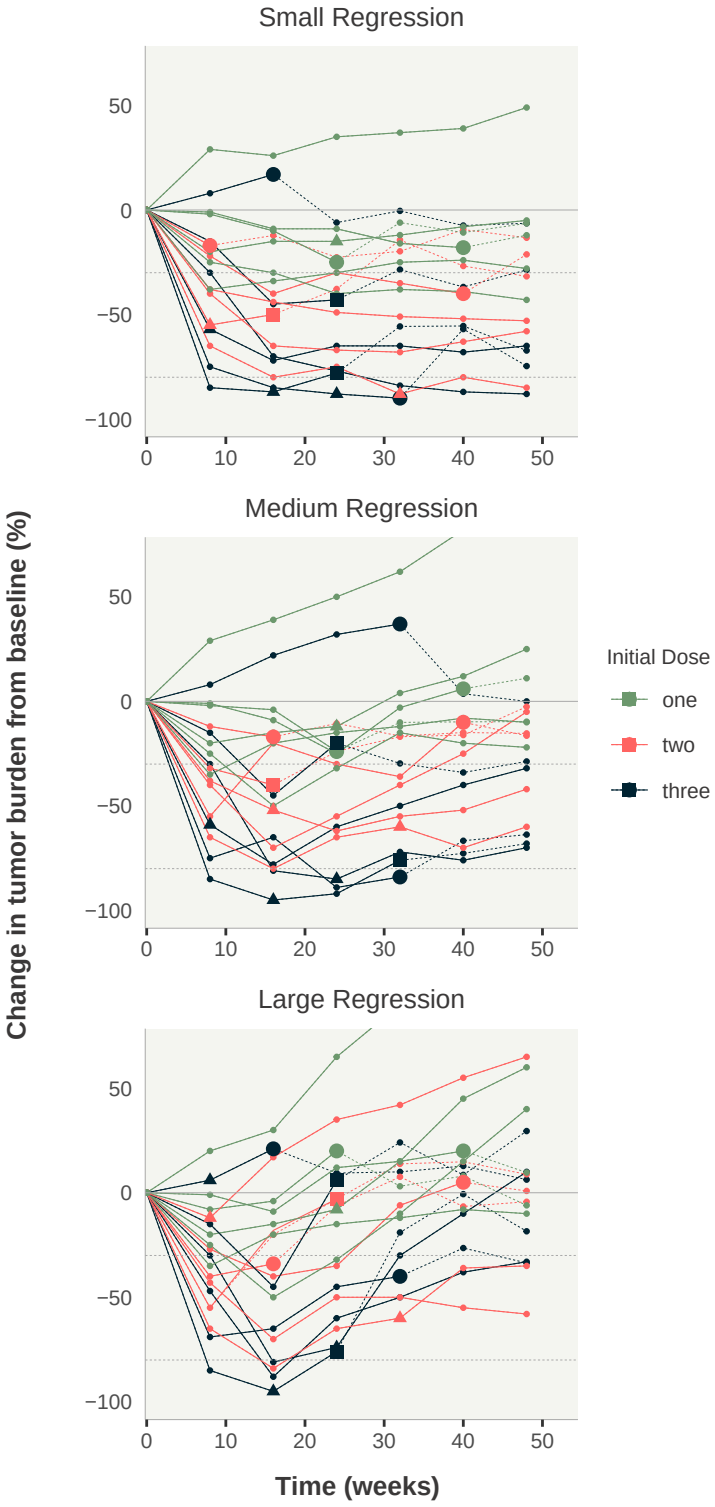

Figure 5.2: Artificial percent tumor change data for three different trajectory situations with two-level imputation ($m = 20$).

Dose	Estimand	Scenario														
		Small Regression					Medium Regression					Large Regression				
		Efficacy					Efficacy					Efficacy				
		prog.	stable	part.	comp.		prog.	stable	part.	comp.		prog.	stable	part.	comp.	
Dose 1	Max. CfB	1	3	2	0	**2**	1	3	2	0	**2**	1	3	2	0	**2**
	Max. CfB (2)	1	3	2	0	**2**	1	4	1	0	**1**	1	4	1	0	**1**
	LOCF	1	4	1	0	**1**	3	3	0	0	**0**	5	1	0	0	**0**
	Imput.	1	4	1	0	**1**	2	4	0	0	**0**	4	2	0	0	**0**
Dose 2	Max. CfB	0	1	4	1	**5**	0	1	4	1	**5**	0	1	4	1	**5**
	Max. CfB (2)	0	1	4	1	**5**	0	1	5	0	**5**	1	1	4	0	**4**
	LOCF	0	1	4	1	**5**	0	3	3	0	**3**	3	0	3	0	**3**
	Imput.	1	1	3	1	**4**	0	4	2	0	**2**	3	1	2	0	**2**
Dose 3	Max. CfB	1	0	2	3	**5**	1	0	2	3	**5**	1	0	2	3	**5**
	Max. CfB (2)	1	0	2	3	**5**	1	1	1	3	**4**	1	1	2	2	**4**
	LOCF	1	0	3	2	**5**	1	1	3	1	**4**	2	1	3	0	**3**
	Imput.	1	1	3	1	**4**	1	1	4	0	**4**	3	1	2	0	**2**

Table 5.6: Responder analysis of artificial data using traditional and proposed estimands.

efficacy again (5/6, 4/6, 3/6). Therefore, the Imputation estimand would recommend no doses to move forward, while the previous estimands would each recommend Dose 2 and Dose 3.

When an obvious U-shape is present, the chosen estimand clearly affects which doses are considered efficacious. Depending on how closely tumor burden correlates to a survival endpoint, this difference could mean preventing non-efficacious (and possibly toxic) doses from being studied in late phase trials, saving time, money, and sub-optimal patient outcomes. Ultimately, though, the desired estimand should be chosen to best address the treatment effect of interest. The Imputation estimand is fundamentally better in this regard than other potential estimands.

An imputation procedure can be effective using a responder outcome, as shown, but would be more appropriate using raw tumor change data or simply raw tumor data. (Tumor change itself is not an ideal response, see Chapter 6.) Regarding the

estimand definition, only the Population Level Summary element would change. An analogous analysis to the responder method would be to set a target of 30% average tumor reduction for a given dose. Table 5.7 shows average percent change in tumor burden for each dose, estimand, and scenario.

Dose	Estimand	Scenario		
		Small	Medium	Large
Dose 1	Max. CfB	-20	-20	-17
	Max. CfB (2)	-14	-9	-7
	LOCF	-12	13	37
	Imput.	-8	16	26
Dose 2	Max. CfB	-54	-54	-50
	Max. CfB (2)	-48	-42	-32
	LOCF	-51	-31	-10
	Imput.	-44	-23	-4
Dose 3	Max. CfB	-62	-64	-62
	Max. CfB (2)	-60	-53	-49
	LOCF	-58	-41	-20
	Imput.	-54	-42	-6

Table 5.7: Average percent tumor burden change-from-baseline of artificial data using traditional and proposed estimands.

Many of the trends from the responder analysis table are consistent, except now shown in more detail. The Max CfB estimand is almost identical for a given dose over regression scenarios, while each estimand is fairly similar for a given dose under the Small Regression scenario. Doses 2 and 3 would be recommended by all estimands in Small Regression. In Medium Regression, LOCF and Imputation are different from the other estimands; moreover, there is a significant difference between each other on Dose 2, where LOCF would recommend (-31), but Imputation would not (-23), based

on sample estimates alone. All estimands would recommend Dose 3 here. In Large Regression, neither LOCF nor Imputation would recommend any of the three doses (a change for LOCF from responder analysis), while the traditional estimands would recommend Dose 2 and Dose 3 as before.

5.3 Dose Selection

Under the current framework, dose is assigned during the trial utilizing toxicity information only. At the end of a trial, the dose x_j^* is the maximum dose with estimated toxicity of probability under both targets π_{DR} and π_{DLT}. Thus, all doses x_j^* or lower are considered safe, if x_j^* exists. Let the set of safe doses be defined as X_-^*.

$$X_-^* = \{x_j \in X : x_j \leq x_j^*\}. \tag{5.3}$$

It may be that there is a plateau or peak in efficacy such that consecutive doses provide similar efficacy. To determine which dose or doses are recommended in order to move forward, the question then becomes whether doses in X_-^* below x_j^* can provide an adequate efficacy profile with significantly lower toxicity. A definition must be given then for what is required of efficacy to be adequate and what is required toxicity to be significantly lower.

Let there be an efficacy threshold π_E (binary outcome) or μ_E (continuous outcome) specified before the trial (if none is specified, $(\pi_E, \mu_E) = 0$) where X_{E_-} is the set of doses in X_{j-}^* such that the estimated efficacy of each dose is greater than or

equal to π_E or μ_E.

$$X_{E_-} = \{x_j \in X_-^* : \hat{\pi}_{E,j} \geq \pi_E\} \tag{5.4}$$

$$X_{E_-} = \{x_j \in X_-^* : \hat{\mu}_{E,j} \geq \mu_E\}. \tag{5.5}$$

If the set X_{E_-} is empty, the drug is considered too toxic or not efficacious enough to continue study. Let the maximum dose in this set be represented as $x_{E_-}^{max}$. If the proposed method is being used in a combined Phase I/II setting with a relatively larger sample size, and a single dose is desired to begin an efficacy registration trial, then the "best dose among safe doses" $x_{E_-}^{max}$ is recommended. This is the Optimal Biological Dose as defined in Chapter 1.

If the proposed method is being used in a strictly Phase I dose-finding setting with a relatively smaller sample size, and a separate Phase II toxicity and efficacy evaluation planned, it is beneficial to recommend multiple doses to move forward (see Chapter 6 for discussion). The following simple algorithm can be used. When X_{E_-} contains one or two doses, then all of X_{E_-} is selected for further study. When X_{E_-} contains three or more doses, define an adequacy proportion parameter a_E and a toxicity proportion parameter b_T, $0 < a_E, b_T \leq 1$. A suggestion would be $a_E = 0.75$ and $a_T = 0.25$. If $x_{E_-}^{max-1}$ and $x_{E_-}^{max-2}$ in X_{E_-} have estimated efficacy within $a_E * \hat{\pi}_{E,max}$ and toxicity lower than $(1 - a_T) * \hat{\pi}_{DR,max}$ and $(1 - a_T) * \hat{\pi}_{DLT,max}$, then $x_{E_-}^{max-1}$ and $x_{E_-}^{max-2}$ are selected for further study without $x_{E_-}^{max}$. The dose x_E^{max-3}, if it exists, is never considered for further study. In other words, using the suggested values, doses $x_{E_-}^{max-1}$ and $x_{E_-}^{max}$ are selected for further study, unless estimated efficacy on $x_{E_-}^{max-1}$ and $x_{E_-}^{max-2}$ is within 75% and estimated toxicity is reduced by more than 25% (Dose Reductions and DLTs) of $x_{E_-}^{max-1}$.

Chapter 6

Discussion

Decisions are guided by limited data setting.

When operating in Phase I and Phase I/II trials, the greatest limitation is sample size. These trials are necessarily "small" data. This fact has guided each modeling decision within this dissertation. The forms of the conditional hazards in the toxicity model can be viewed as under-parameterized. This decision was made, though, for the same reasons the original CRM is a one-parameter model. Estimating additional parameters with such limited data comes at a great variance cost. The current specifications are an attempt to retain important assumptions about the data while minimizing the number of parameters.

The Bayesian framework for estimation is adopted for two reasons, common to the literature. With toxic drugs (even if less toxic than chemotherapies), it is ethical to update the model throughout the trial, so that new patients are not enrolled at dangerous doses based on the data at hand. The Bayesian framework allows for this updating without the need for multiple comparison adjustments to control some Type I error rate. Similarly, data already collected can be used effectively to inform the dose that a newly enrolled patient receives.

Sample size of $n = 40$ is recommended for toxicity evaluation.

Based on simulation results, a sample size of $n = 40$ provides good operating characteristics and selection performance based on toxicity estimation alone. Con-

versely, sample sizes of $n = 10$ and $n = 20$ are quite noisy and biased, and should not be considered for a robust evaluation of toxicity. A conservative sample size, to better guarantee accurate estimates and selection, is $n = 60$. Beyond $n = 60$, there are negligible improvements in PCS, Overdose percentage, accuracy index, and bias and variance of toxicity probability estimates. Somewhere between $n = 40$ and $n = 60$ should suit efficacy evaluation as well, where $\sim$5 doses are studied. Sample size for efficacy is not studied in this dissertation. To design a proper simulation study would require more information about real trial efficacy.

Responder analysis and change-from-baseline response are flawed.

Regarding efficacy, two aspects of the outcome variable should be discussed further: responder analysis and percent change-from-baseline transformation. A RECIST-type responder analysis is a discretization of a continuous outcome into four, and ultimately two, categories. From Diniz, Tighiouart, and Rogatko 2019, "[d]espite...an extensive statistical literature showing that discretizing continuous variables results in substantial loss of information, categorization of continuous variables has been a common practice in clinical research and in cancer dose finding (Phase I) clinical trials." The context of Diniz's statement was set in the discretization of dose sets, but this general statement is equally applicable to the context of measurable endpoints. Much work has been accomplished in various contexts attempting to discretize continuous variables optimally. However, the fact remains that discretizing a continuous variable leads to a loss of information. Modeling longitudinal tumor burden directly is preferred in order to utilize the maximum information available, when limited information is available to begin with, instead of an arbitrary dichomtomization.

Additionally, although change-from-baseline outcomes are ubiquitous in clinical

trials, a strong argument can be made that they are inappropriate in parallel group studies, of which a dose-finding trial is one, especially percent change-from-baseline. The purpose of a parallel-group study is not to compare a patient with themselves at baseline, but to compare the groups. Per Harrell 2022, "Within-patient change is affected strongly by regression to the mean and measurement error. When the baseline value is one of the patient inclusion/exclusion criteria, the only meaningful change score requires one to have a second baseline measurement post patient qualification to cancel out much of the regression to the mean effect. It is the second baseline that would be subtracted from the follow-up measurement." Additionally, the baseline value must be linearly related to the post value to allow for valid analysis.

Percent change-from-baseline compounds these issues by creating an asymmetric measure where subtracting itself does not make sense. Suppose a patient started with 5cm tumor burden and increased to 10cm, while another patient started at 10cm tumor burden and decreased to 5cm. This effect should be zero; however, percent change-from-baseline is 100% and -50%, averaging out to +25% change. The solution would be to analyze raw tumor burden value. As mentioned in Chapter 5, the proposed imputation procedure could be used with such a continuous response. The illustration focuses on a responder analysis and percent change-from-baseline outcome, despite their faults, because of ubiquity, that a researcher may see immediate value in applying the current methods.

Multiple imputation with a mixed model allows necessary flexibility.

Repeated Measures ANCOVA compares responses of groups when individuals have repeated measures, but there are conflicts with the assumptions of this model and the current efficacy setting. Time of visit is often not fixed and visits are unbalanced, as evidenced in the example trials, in which case a mixed model treating

time as continuous is necessary. Subsequently, the effect of time can be modeled as nonlinear. Also, the goal of evaluating efficacy is likely identifying which dose levels exceed a separate, prespecified target value, rather than identifying which dose levels are significantly different from each other. Imputing continuous values better allows for this comparison, whether the target value is an average tumor burden reduction or a categorized response rate.

Selecting more than one dose is encouraged.

As suggested in Section 2.5, more than one dose should be evaluated until at least the start of a Phase III registration trial. Targeted therapies are expected to have long-term and cumulative toxicities; and while proper design of the trial can create an efficient environment, more time and information is necessary for proper evaluation (Janne et al. 2016). Targeted therapies often exhibit a delicate balance of benefits and risks; and with the greater possibility of an efficacy plateau, evaluating toxicity and efficacy more thoroughly over a range of doses is beneficial (Bullock, Rahman, and Liu 2016). The proposed toxicity and efficacy methods can be utilized in a small Phase I dose-finding trial that suggests multiple doses to move forward. Or, it can be utilized in a more expansive Phase I/II setting that evaluates multiple doses throughout, choosing a single dose at the end.

TITECRM-MC significantly contrasts with the proposed toxicity method.

As TITECRM-MC is the only other method in the literature that addresses dose reductions by incorporating both graded toxicity and partial information, the important differences with the proposed toxicity method should be highlighted. The main difference is that toxicity in TITECRM-MC is ordinal. This means a patient can be recorded as a Moderate Toxicity (Dose Reduction) or a DLT, but not both.

If a patient has a Dose Reduction and a subsequent DLT, the model considers that patient only as a DLT. Conversely, the proposed method would consider one Dose Reduction and one DLT for such a situation.

An ordinal outcome also means that all toxicity is associated with initial dose. In the situation described previously, the DLT is associated with the patient's initial dose, even though the patient is on a lower dose between the time of Dose Reduction and DLT. The reduced dose in never included in the model. It is likely in trial protocol that an interruption period is applied after a Dose Reduction to allow the initial toxicity to resolve. This means that a subsequent DLT would be a combination of the cumulative effects of the initial dose and the reduced dose, but more appropriately associated with the reduced dose. The proposed method is fit to handle both Dose Reduction and Dose Interruption situations. TITECRM-MC is more appropriately designed for inclusion of Dose Interruptions only, during which a patient stays on their initial dose.

If the likelihood weight parameter of TITECRM-MC is linearly constant, as is common, this would also be unfavorable, as the probability of toxicity will most likely be decreasing over the time window. This is evidenced in the simulations, where the proposed method performed better on data generated by decreasing hazard rates. The proposed method does pay a price by estimating four parameters instead of two, but seems to benefit.

6.1 Novel Contributions

Listed below are what are assumed to be the novel contributions in this dissertation.

- The application of a semi-competing risks framework to dose-finding trials.

- Constant-Skeleton and Weibull-Skeleton conditional hazard formulations.
- A toxicity model that can incorporate both Dose Interruptions and Dose Reductions separately.
- A method that explicitly controls the rate of Dose Reductions.
- The development of the likelihood associated with the five-transition model.
- The development of the transition probabilities associated with the five-transition model.
- An original derivation of the likelihood and transition probabilities associated with the three-transition or illness-death model.
- A simulation study of the Constant-Skeleton and Weibull-Skeleton illness-death models in a realistic setting compared with other designs used in current clinical settings.
- Asymptotic proof that the MUTD dose selection mechanism using derived transition probabilities will choose the desired dose.
- Recommendation of appropriate sample size when using the proposed toxicity method based on simulation.
- Application of two-level multiple imputation to handle missing tumor reduction data resulting from DLTs while incorporating Dose Reductions.
- A novel efficacy estimand for binary and continuous outcomes in tumor burden context.
- Use of the above methods in order to make dosing decisions in an oncological setting.

6.2 Limitations and Future Work

Nowadays, targeted agents and immunotherapies are often given in combination with each other or in combination with traditional therapies. Drug combination strategies are not discussed in this dissertation, but could be explored to make the proposed framework more robust for application. A reasonable strategy to handle multiple drugs is to specify a set of "simple orderings" from the known "partial orderings" induced by a two-dimensional or n-dimensional grid of doses, as detailed in Wages, Conaway, and O'Quigley 2011. Then, when dosing patients, the observed data can be used to first choose the most likely simple ordering of dose combinations before choosing the best dose combination within that simple ordering.

Moreover, a common characteristic of initial dose-finding trials is a stopping rule, where a trial can be terminated early for success or failure based on overwhelming toxicity or efficacy data. With stopping rules also come multiple comparison considerations. Using the data to make inferences at multiple time points, each with its own error rate, requires adjustments to each individual inference to control an overall decision-making error rate. Stopping rules and related multiple comparison considerations are not discussed in this dissertation. Similarly, in the simulations, never was it the case that zero doses were chosen. Instead, if all doses were estimated to be too toxic, dose level 1 was chosen. It may be interesting to see how often the proposed method rejects all doses as too toxic when it should do such, as in simulation Scenario 7.

The most complex toxicity model considered was the Weibull-Skeleton model, a four parameter model. More complexity is unlikely to be beneficial. But because the Weibull-Skeleton became the preferred model, a model with an extra parameter could be considered to explore additional complexity. This extra parameter could add

flexibility to dose scaling or represent patient heterogeneity. Or rather, the Weibull-Skeleton model could have β removed, leaving dose scaling to the skeleton itself, trying fit non-constant hazards with a three-parameter model.

Obviously, more priors and more data generating scenarios could be considered in simulation. If feedback identifies additional scenarios that were not considered, further simulations could be added. One situation that may be of interest in the context of the illness-death model is when a patient may have a Dose Interruption or a Dose Reduction. Similarly, simulations could be conducted on the five-transition model, assuming there could be two Dose Reductions, or that each patient could have a Dose Interruption and a Dose Reduction. Also, the dosing decision rules could be further modified in order to find an optimal balance of aggressiveness in dose escalation and restriction. The dosing decision rule used in the TITECRM-MC (Lee et al. 2019) is different from any considered in the proposed method, and may present an interesting balance in cautiousness.

With regard to the Multiple Imputation model for efficacy, the optimal set of covariates for imputation of the response may be different depending on the specific drug and population setting. For example, it may be beneficial to structurally differentiate between responders and non-responders. From the real trial efficacy shown in this dissertation, there is reliably a set of patients that immediately respond to the drug, seeing their tumor burden decrease in the first follow-up visits, and a set a patients that do not respond to the drug, seeing their tumor burden stay constant or only increase from the start. Two populations thus emerge. Values from initial visits could be used to define these populations and subsequently included as a covariate in the imputation model. This should help generate more realistic imputations for those whom experience initial increase in tumor burden, as in the current illustration there is regression to the mean which may not be most realistic. Additionally, the

multinomial model used to impute dose level was not a two-level model because such a model is not readily available to use for imputation. The focus of the current efficacy work is on the estimand and missing data assumptions and not the technical modeling aspects, yet in the future, more work could be done to adapt or develop a multinomial imputation model for this purpose.

Bibliography

O'Quigley, John, Margaret Pepe, and Lloyd Fisher (1990). "Continual reassessment method: a practical design for phase 1 clinical trials in cancer". In: *Biometrics*, pp. 33–48.

Wei, Lee-Jen (1992). "The accelerated failure time model: a useful alternative to the Cox regression model in survival analysis". In: *Statistics in Medicine* 11.14-15, pp. 1871–1879.

O'Quigley, John and Larry Z. Shen (1996). "Continual reassessment method: a likelihood approach". In: *Biometrics*, pp. 673–684.

Hougaard, Philip (1999). "Multi-state Models: A Review". In: *Lifetime Data Analysis* 5, pp. 239–264.

Cheung, Ying Kuen and Rick Chappell (2000). "Sequential designs for phase I clinical trials with late-onset toxicities". In: *Biometrics* 56.4, pp. 1177–1182.

Schafer, Joseph L. and Recai M. Yucel (2002). "Computational strategies for multivariate linear mixed-effects models with missing values". In: *Journal of Computational and Graphical Statistics* 11.2, pp. 437–457.

Gadbury, GL, CS Coffey, and DB Allison (2003). "Modern statistical methods for handling missing repeated measurements in obesity trial data: beyond LOCF". In: *Obesity Reviews* 4.3, pp. 175–184.

Thall, Peter F. et al. (2003). "Dose-finding with two agents in phase I oncology trials". In: *Biometrics* 59.3, pp. 487–496.

Rubin, Donald B. (2004). *Multiple Imputation for Nonresponse in Surveys*. Vol. 81. John Wiley & Sons.

Bekele, B. Nebiyou and Yu Shen (2005). "A Bayesian approach to jointly modeling toxicity and biomarker expression in a phase I/II dose-finding trial". In: *Biometrics* 61.2, pp. 343–354.

Johnson, John R. et al. (2005). "Approval summary for erlotinib for treatment of patients with locally advanced or metastatic non–small cell lung cancer after failure of at least one prior chemotherapy regimen". In: *Clinical Cancer Research* 11.18, pp. 6414–6421.

Philip, Philip A. et al. (2005). "Phase II study of Erlotinib (OSI-774) in patients with advanced hepatocellular cancer". In: *Journal of Clinical Oncology* 23.27, pp. 6657–6663.

Wang, Kai and Anastasia Ivanova (2005). "Two-dimensional dose finding in discrete dose space". In: *Biometrics* 61.1, pp. 217–222.

Braun, Thomas M. (2006). "Generalizing the TITE-CRM to adapt for early-and late-onset toxicities". In: *Statistics in Medicine* 25.12, pp. 2071–2083.

Putter, Hein, Marta Fiocco, and Ronald B. Geskus (2007). "Tutorial in biostatistics: competing risks and multi-state models". In: *Statistics in Medicine* 26.11, pp. 2389–2430.

Streiner, David L (2008). "Missing data and the trouble with LOCF". In: *BMJ Ment Health* 11.1, pp. 3–5.

Eisenhauer, Elizabeth A. et al. (2009). "New response evaluation criteria in solid tumours: Revised RECIST guideline (version 1.1)". In: *European Journal of Cancer* 45.2, pp. 228–247.

Yin, Guosheng and Ying Yuan (2009). "Bayesian model averaging continual reassessment method in phase I clinical trials". In: *Journal of the American Statistical Association* 104.487, pp. 954–968.

Braun, Thomas M. and Shufan Wang (2010). "A hierarchical Bayesian design for phase I trials of novel combinations of cancer therapeutic agents". In: *Biometrics* 66.3, pp. 805–812.

Hernán, Miguel A. (2010). "The hazards of hazard ratios". In: *Epidemiology (Cambridge, Mass.)* 21.1, p. 13.

Rodriguez, German (2010). *Multivariate Survival Models*.

Xu, Jinfeng, John D. Kalbfleisch, and Beechoo Tai (2010). "Statistical analysis of illness–death processes and semicompeting risks data". In: *Biometrics* 66.3, pp. 716–725.

Cheung, Ying Kuen (2011). *Dose Finding by the Continual Reassessment Method*. CRC Press.

Lee, Bin Cheng, and Ying Kuen Cheung (2011). "Continual reassessment method with multiple toxicity constraints". In: *Biostatistics* 12.2, pp. 386–398.

Mauguen, A., M. C. Le Deley, and S. Zohar (2011). "Dose-finding approach for dose escalation with overdose control considering incomplete observations". In: *Statistics in Medicine* 30.13, pp. 1584–1594.

Van Buuren, Stef (2011). "Multiple imputation of multilevel data". In: *The Handbook of Advanced Multilevel Analysis*. Ed. by J.J. Hox and J.K. Roberts. Milton Park, UK: Routledge. Chap. 10, pp. 173–196.

Van Buuren, Stef and Karin Groothuis-Oudshoorn (2011). "mice: Multivariate Imputation by Chained Equations in R". In: *Journal of Statistical Software* 45.3, pp. 1–67.

Wages, Nolan A., Mark R. Conaway, and John O'Quigley (2011). "Dose-finding design for multi-drug combinations". In: *Clinical Trials* 8.4, pp. 380–389.

Yuan, Ying and Guosheng Yin (2011). "Robust EM continual reassessment method in oncology dose finding". In: *Journal of the American Statistical Association* 106.495, pp. 818–831.

Lee, Shing M., D. L. Hershman, et al. (2012). "Toxicity burden score: a novel approach to summarize multiple toxic effects". In: *Annals of Oncology* 23.2, pp. 537–541.

Meter, Emily M. Van, Elizabeth Garrett-Mayer, and Dipankar Bandyopadhyay (2012). "Dose-finding clinical trial design for ordinal toxicity grades using the continuation ratio model: an extension of the continual reassessment method". In: *Clinical Trials* 9.3, pp. 303–313.

Gelman, Andrew et al. (2013). *Bayesian Data Analysis*. CRC press.

Hamid, Omid et al. (2013). "Safety and tumor responses with lambrolizumab (anti–PD-1) in melanoma". In: *New England Journal of Medicine* 369.2, pp. 134–144.

Liu, Suyu and Jing Ning (2013). "A Bayesian dose-finding design for drug combination trials with delayed toxicities". In: *Bayesian Analysis* 8.3, p. 703.

Liu, Suyu, Guosheng Yin, and Ying Yuan (2013). "Bayesian data augmentation dose finding with continual reassessment method and delayed toxicity". In: *The Annals of Applied Statistics* 7.4, p. 1837.

Wolchok, Jedd D. et al. (2013). "Nivolumab plus ipilimumab in advanced melanoma". In: *New England Journal of Medicine* 369, pp. 122–133.

Postel-Vinay, S. et al. (2014). "Towards new methods for the determination of dose limiting toxicities and the assessment of the recommended dose for further studies of molecularly targeted agents–dose-Limiting Toxicity and Toxicity Assessment Recommendation Group for Early Trials of Targeted therapies, an European Organisation for Research and Treatment of Cancer-led study". In: *European Journal of Cancer* 50.12, pp. 2040–2049.

Jolani, Shahab et al. (2015). "Imputation of systematically missing predictors in an individual participant data meta-analysis: a generalized approach using MICE". In: *Statistics in Medicine* 34.11, pp. 1841–1863.

Lee, Kyu Ha, Sebastien Haneuse, et al. (2015). "Bayesian semi-parametric analysis of semi-competing risks data: investigating hospital readmission after a pancreatic cancer diagnosis". In: *Applied Statistics* 64.2, p. 253.

Paoletti, Xavier et al. (2015). "Dose finding with longitudinal data: simpler models, richer outcomes". In: *Statistics in Medicine* 34.22, pp. 2983–2998.

Uno, Hajime et al. (2015). "Alternatives to hazard ratios for comparing the efficacy or safety of therapies in noninferiority studies". In: *Annals of Internal Medicine* 163.2, pp. 127–134.

Wages, Nolan A., Craig L Slingluff Jr, and Gina Petroni (2015). "A phase I/II adaptive design to determine the optimal treatment regimen from a set of combination immunotherapies in high-risk melanoma". In: *Contemporary Clinical Trials* 41, pp. 172–179.

Wages, Nolan A. and Christopher Tait (2015). "Seamless phase I/II adaptive design for oncology trials of molecularly targeted agents". In: *Journal of Biopharmaceutical Statistics* 25.5, pp. 903–920.

Bullock, Julie M., Atiqur Rahman, and Qi Liu (2016). "Lessons learned: dose selection of small molecule–targeted oncology drugs". In: *Clinical Cancer Research* 22.11, pp. 2630–2638.

Janne, Pasi A. et al. (2016). "Dose finding of small-molecule oncology drugs: optimization throughout the development life cycle". In: *Clinical Cancer Research* 22.11, pp. 2613–2617.

Lachin, John M (2016). "Fallacies of last observation carried forward analyses". In: *Clinical trials* 13.2, pp. 161–168.

Lin, Zhoumeng et al. (2016). "A computational framework for interspecies pharmacokinetics, exposure and toxicity assessment of gold nanoparticles". In: *Nanomedicine* 11.2, pp. 107–119.

Roda, Desamparados, Begona Jimenez, and Udai Banerji (2016). "Are doses and schedules of small-molecule targeted anticancer drugs recommended by phase I studies realistic?" In: *Clinical Cancer Research* 22.9, pp. 2127–2132.

Sachs, Jeffrey R. et al. (2016). "Optimal dosing for targeted therapies in oncology: drug development cases leading by example". In: *Clinical Cancer Research* 22.6, pp. 1318–1324.

Seiwert, Tanguy Y. et al. (2016). "Safety and clinical activity of pembrolizumab for treatment of recurrent or metastatic squamous cell carcinoma of the head and neck (KEYNOTE-012): an open-label, multicentre, phase 1b trial". In: *The Lancet Oncology* 17.7, pp. 956–965.

Todd, John A. et al. (2016). "Regulatory T cell responses in participants with type 1 diabetes after a single dose of interleukin-2: a non-randomised, open label, adaptive dose-finding trial". In: *PLOS Medicine* 13.10.

Emens, Leisha A. et al. (2017). "Report on the Third FDA–AACR Oncology Dose-Finding Workshop". In: *Cancer Immunology Research* 5.12, pp. 1058–1061.

ICH, ICH (2017). "E9 (R1): Addendum to Statistical Principles for Clinical Trials on Choosing Appropriate Estimands and Defining Sensitivity Analyses in Clinical Trials". In: *International Conference on Harmonization*.

Lee, Kyu Ha, Virginie Rondeau, and Sebastien Haneuse (2017). "Accelerated failure time models for semi-competing risks data in the presence of complex censoring". In: *Biometrics* 73.4, pp. 1401–1412.

Yoh, Kiyotaka et al. (2017). "Vandetanib in patients with previously treated RET-rearranged advanced non-small-cell lung cancer (LURET): an open-label, multicentre phase 2 trial". In: *The Lancet Respiratory Medicine* 5.1, pp. 42–50.

Hirakawa, Akihiro, Kan Yonemori, et al. (2018). "Potential utility of a longitudinal relative dose intensity of molecularly targeted agents in phase 1 dose-finding trials". In: *Cancer Science* 109.1, pp. 207–214.

Diniz, Marcio Augusto, Mourad Tighiouart, and Andre Rogatko (2019). "Comparison between continuous and discrete doses for model based designs in cancer dose finding". In: *PLOS ONE* 14.1.

Du, Yu et al. (2019). "An adaptive multi-stage phase I dose-finding design incorporating continuous efficacy and toxicity data from multiple treatment cycles". In: *Journal of Biopharmaceutical Statistics* 29.2, pp. 271–286.

Hirakawa, Akihiro, Yuichi Tanaka, and Shuhei Kaneko (2019). "Pragmatic dose-escalation methods incorporating relative dose intensity assessment for molecularly targeted agents in phase I trials". In: *Contemporary Clinical Trials Communications* 16.

Ivanova, Anastasia and Se Hee Kim (2019). "Dose finding for continuous and ordinal outcomes with a monotone objective function: a unified approach". In: *Biometrics* 65, pp. 307–315.

Lee, Shing M., Moreno Ursino, et al. (2019). "Dose-finding designs for cumulative toxicities using multiple constraints". In: *Biostatistics* 20.1, pp. 17–29.

Gogtay, Nithya Jaideep, Priya Ranganathan, and Rakesh Aggarwal (2021). "Understanding estimands". In: *Perspectives in Clinical Research* 12.2, pp. 106–112.

Harrell Jr., Frank E. (2022). *Biostatistics for Biomedical Research*. URL: https://hbiostat.org/bbr/.

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

Stan Development Team (2022a). *RStan: the R interface to Stan*. R package version 2.21.8. URL: http://mc-stan.org/.

— (2022b). *Stan Modeling Language Users Guide and Reference Manual*. Version 2.21.0. URL: http://mc-stan.org/.

Common Terminology Criteria for Adverse Events (CTCAE) (Feb. 2023). URL: https://ctep.cancer.gov/protocolDevelopment/electronic_applications/ctc.htm.

Appendices

Appendix A

Figures and Tables

A.1 Tumor burden from real trials

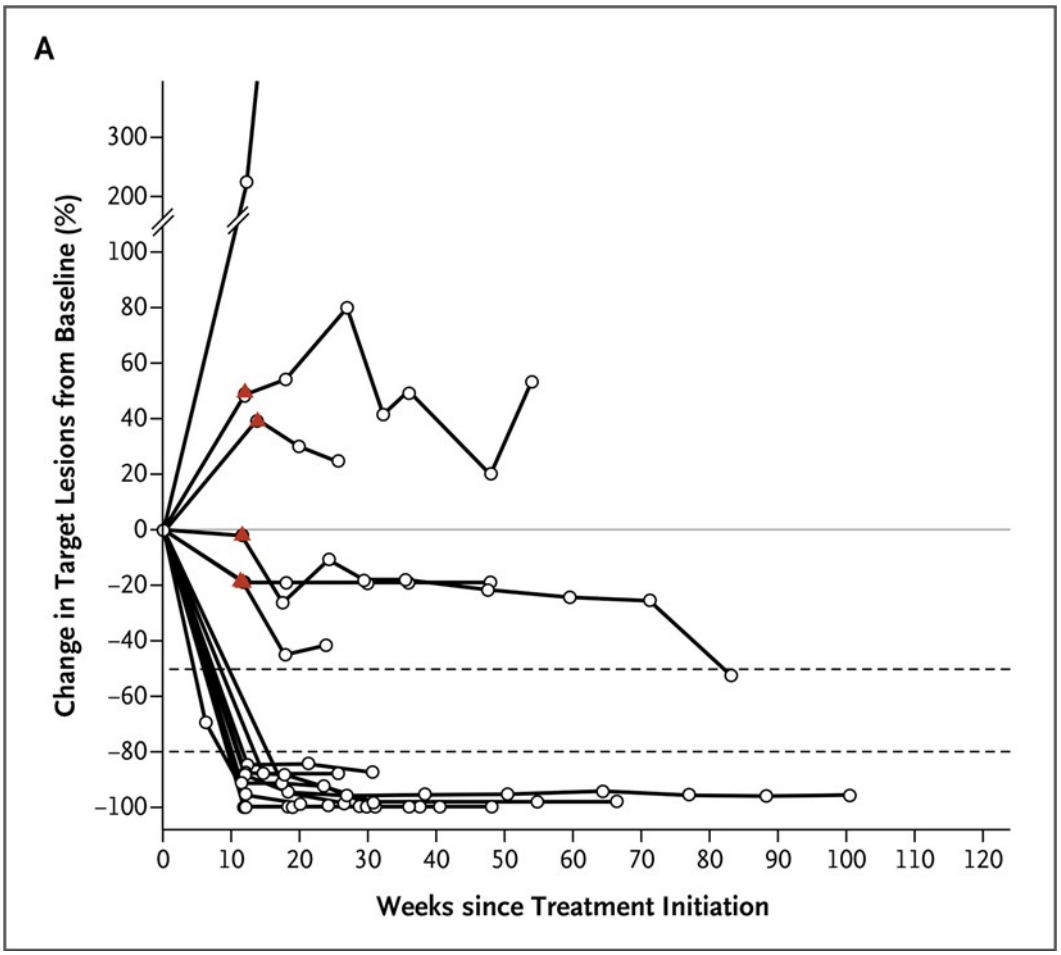

Figure A.1: Percent change from baseline in tumor burden for nivolumab plus ipilimumab in patients with advanced melanoma from Wolchok et al. 2013.

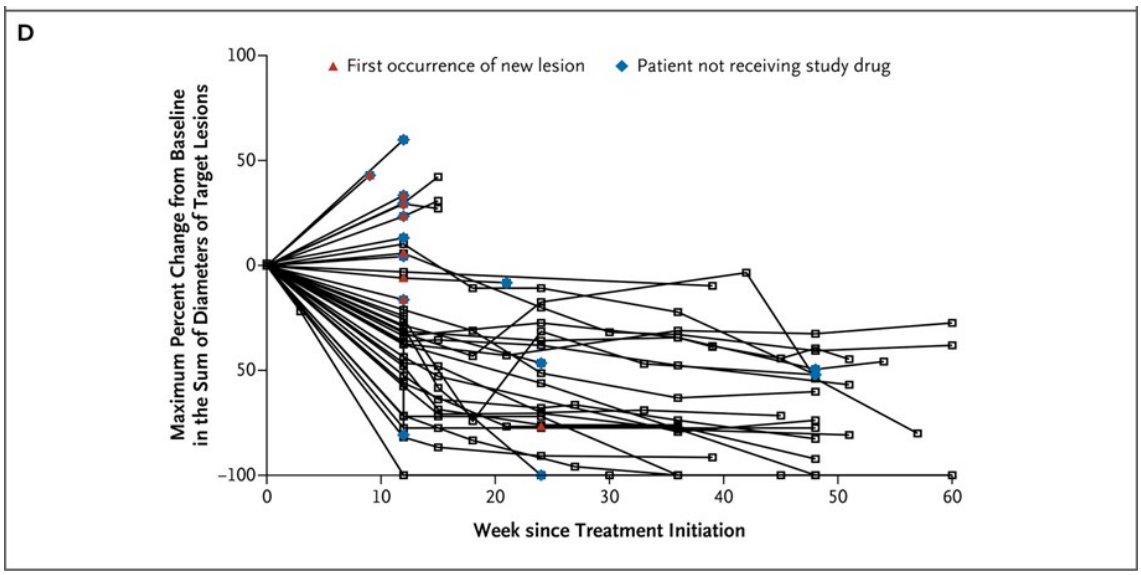

Figure A.2: Percent change from baseline in tumor burden for lambrolizumab in patients with advanced melanoma from Hamid et al. 2013.

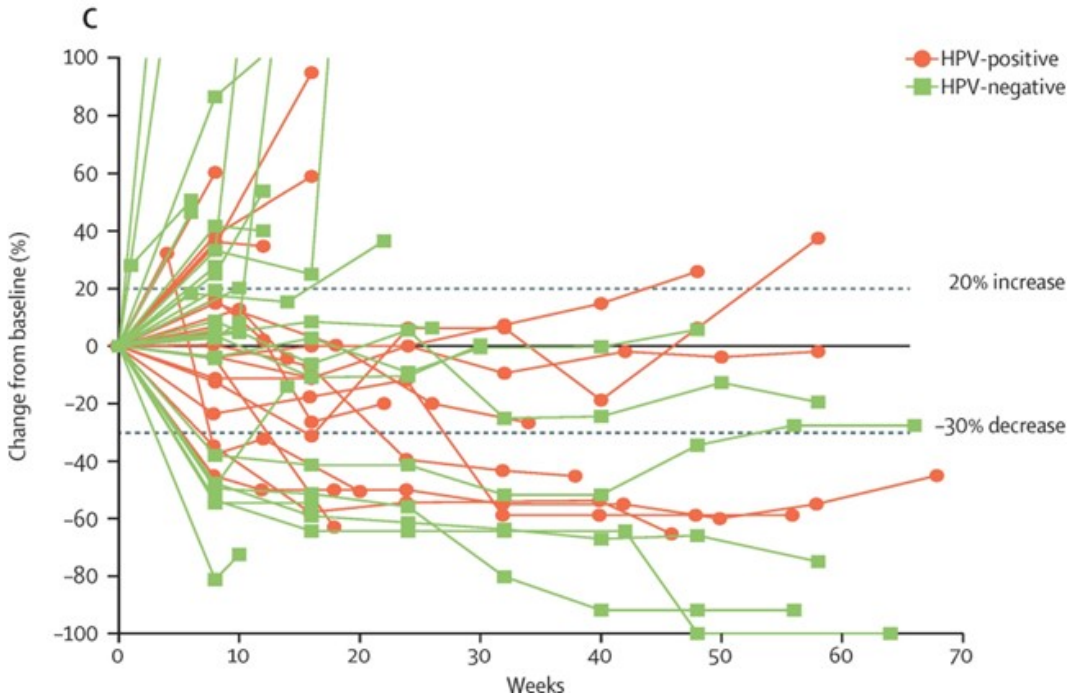

Figure A.3: Percent change from baseline in tumor burden for pembrolizumab in patients with recurrent or metastatic squamous cell carcinoma of the head and neck from Seiwert et al. 2016.

A.2 Simulation supplement

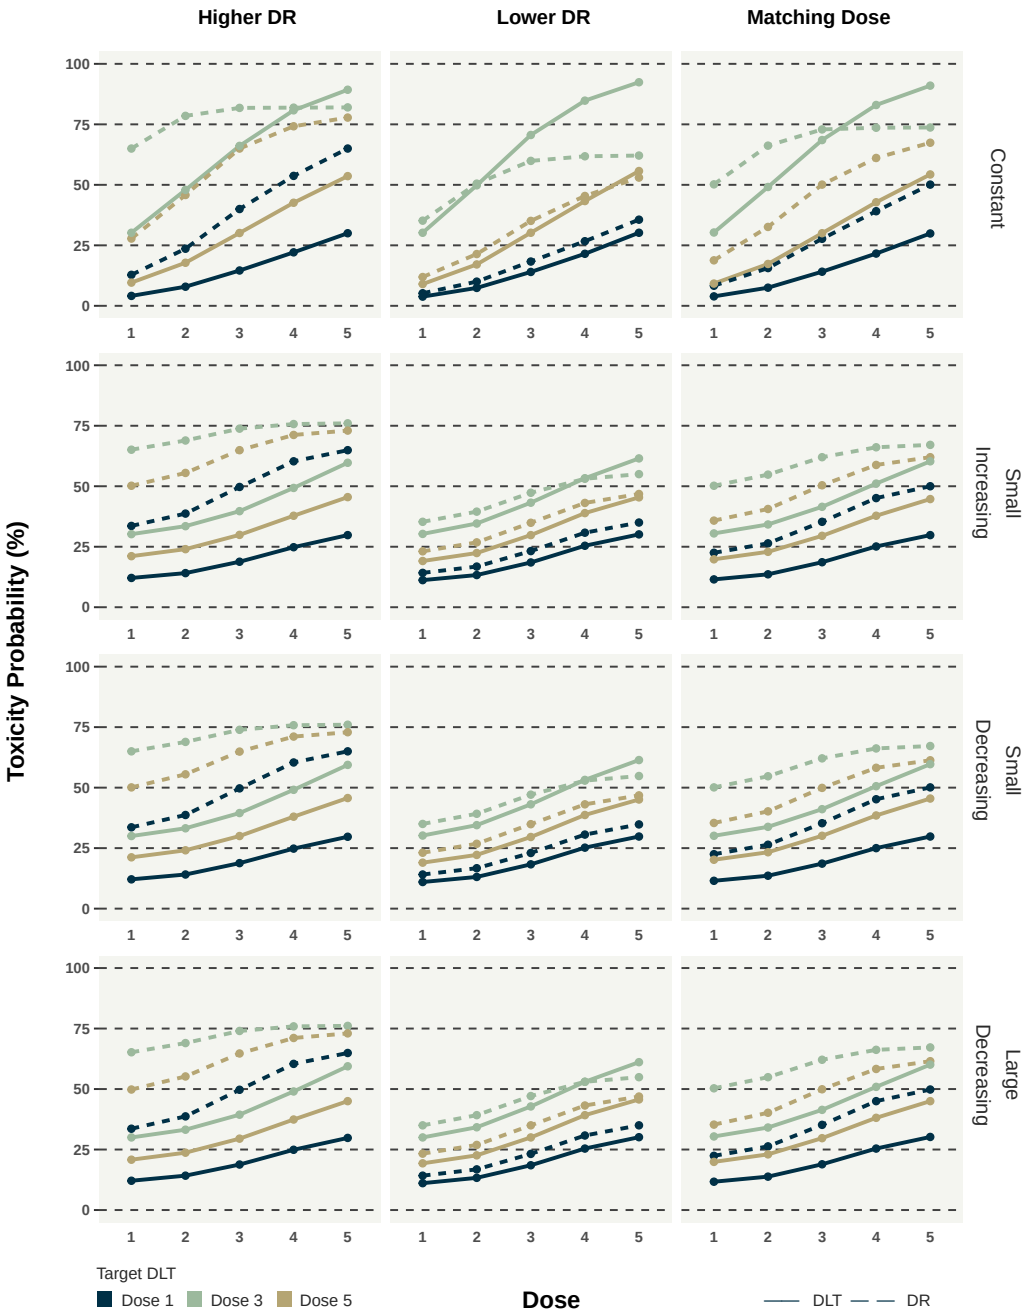

Figure A.4: Examples of dose-toxicity curves for all data generating hazards ($\pi_{DR} = 0.5, \pi_{DLT} = 0.3$)

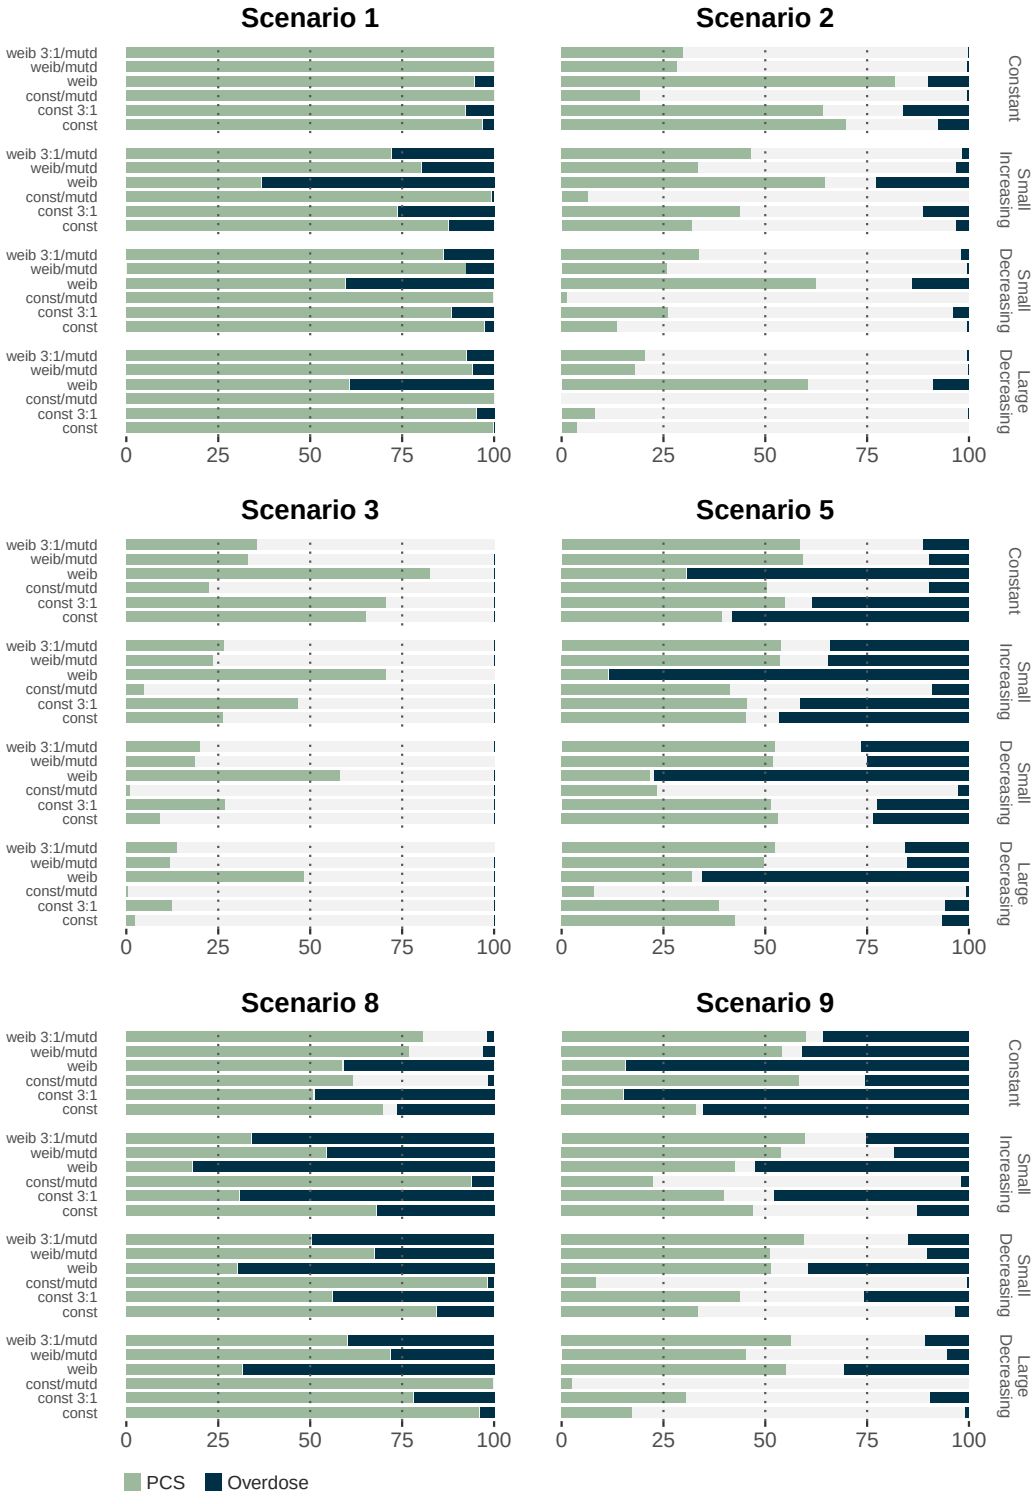

Figure A.5: Percent Correct Selection (green) and overdose (navy) for each candidate model combination over all data generating hazards (n=60).

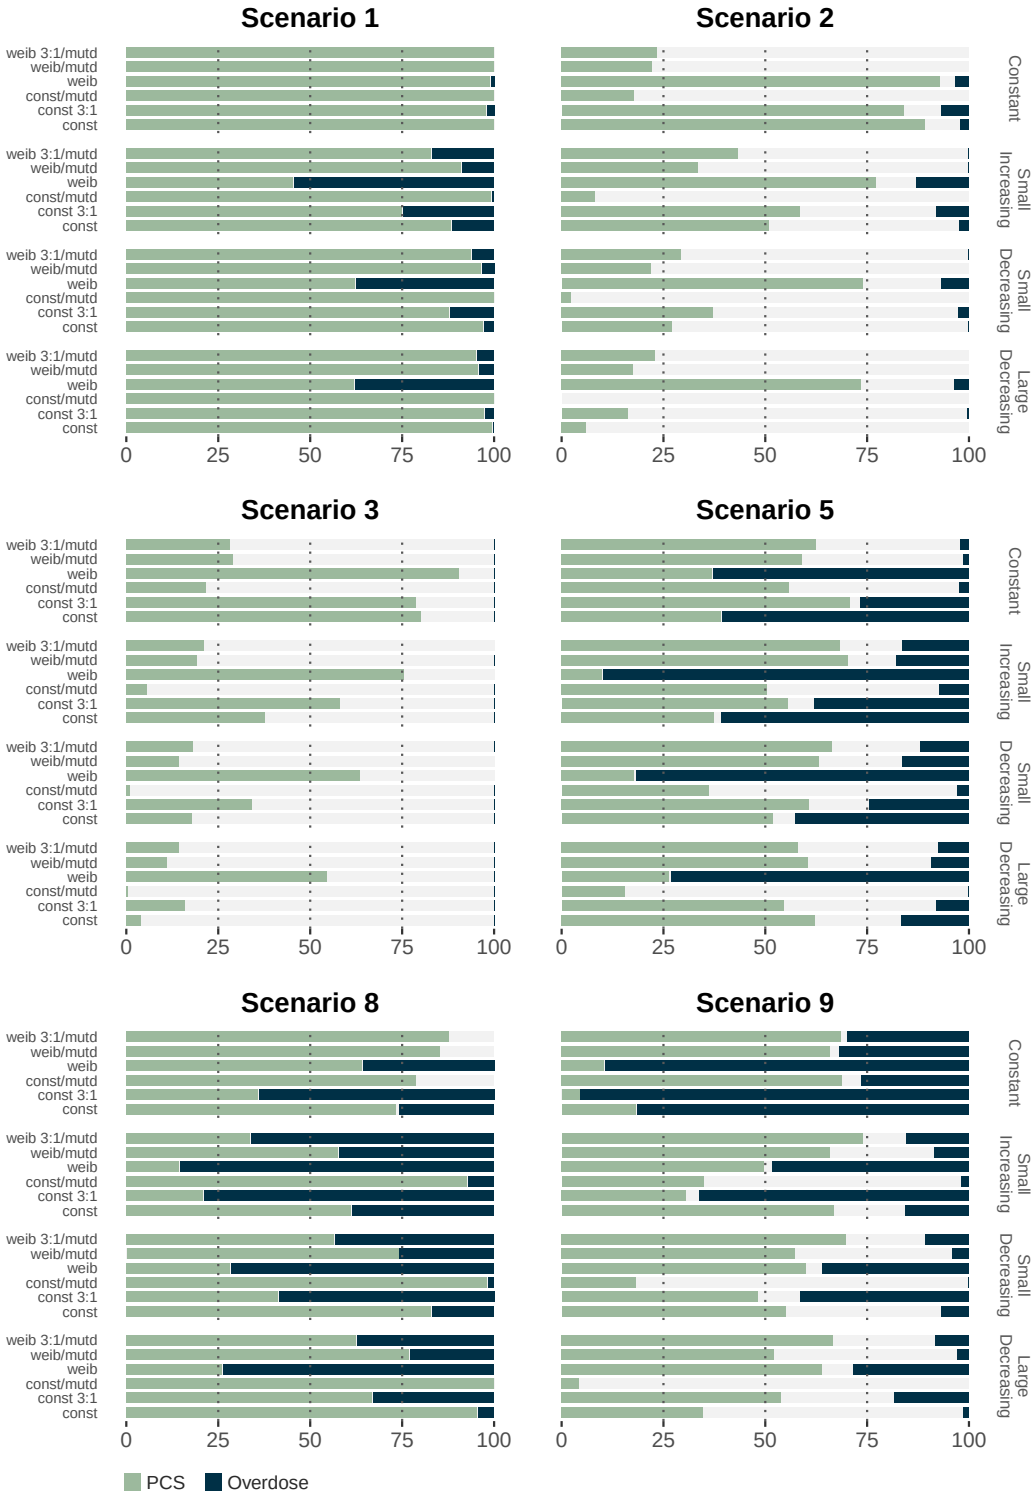

Figure A.6: Percent Correct Selection (green) and overdose (navy) for each candidate model combination over all data generating hazards (n=100).

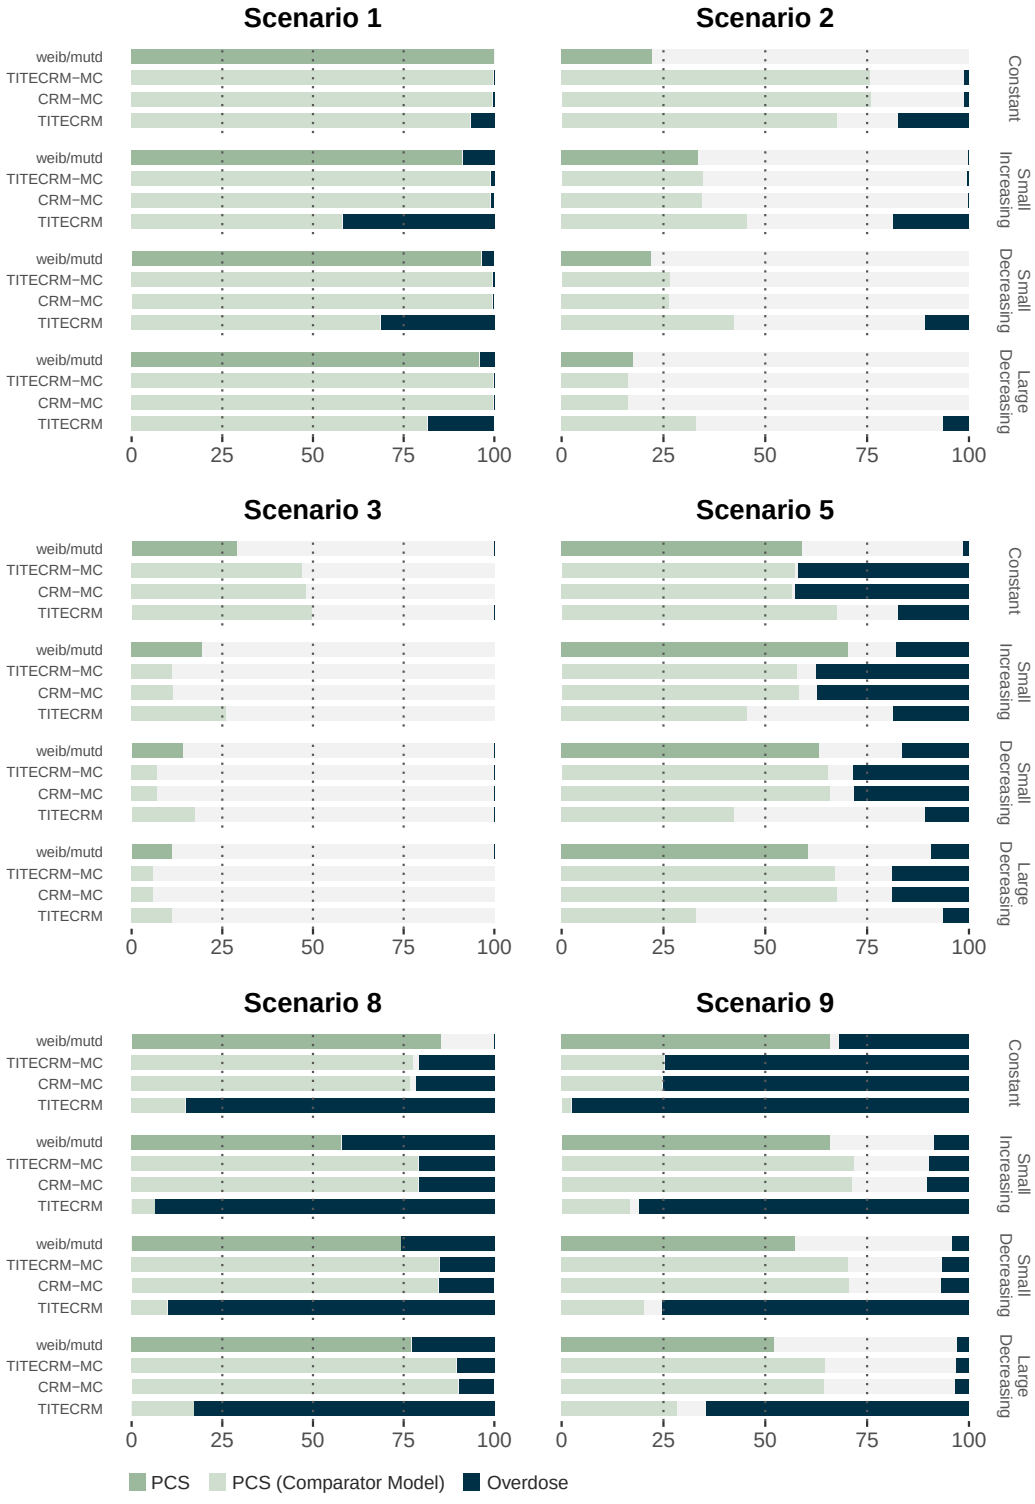

Figure A.7: Percent Correct Selection (green/light green) and Overdose (navy) for Weibull-Skeleton + MUTD candidate model and comparator models ($n = 100$).

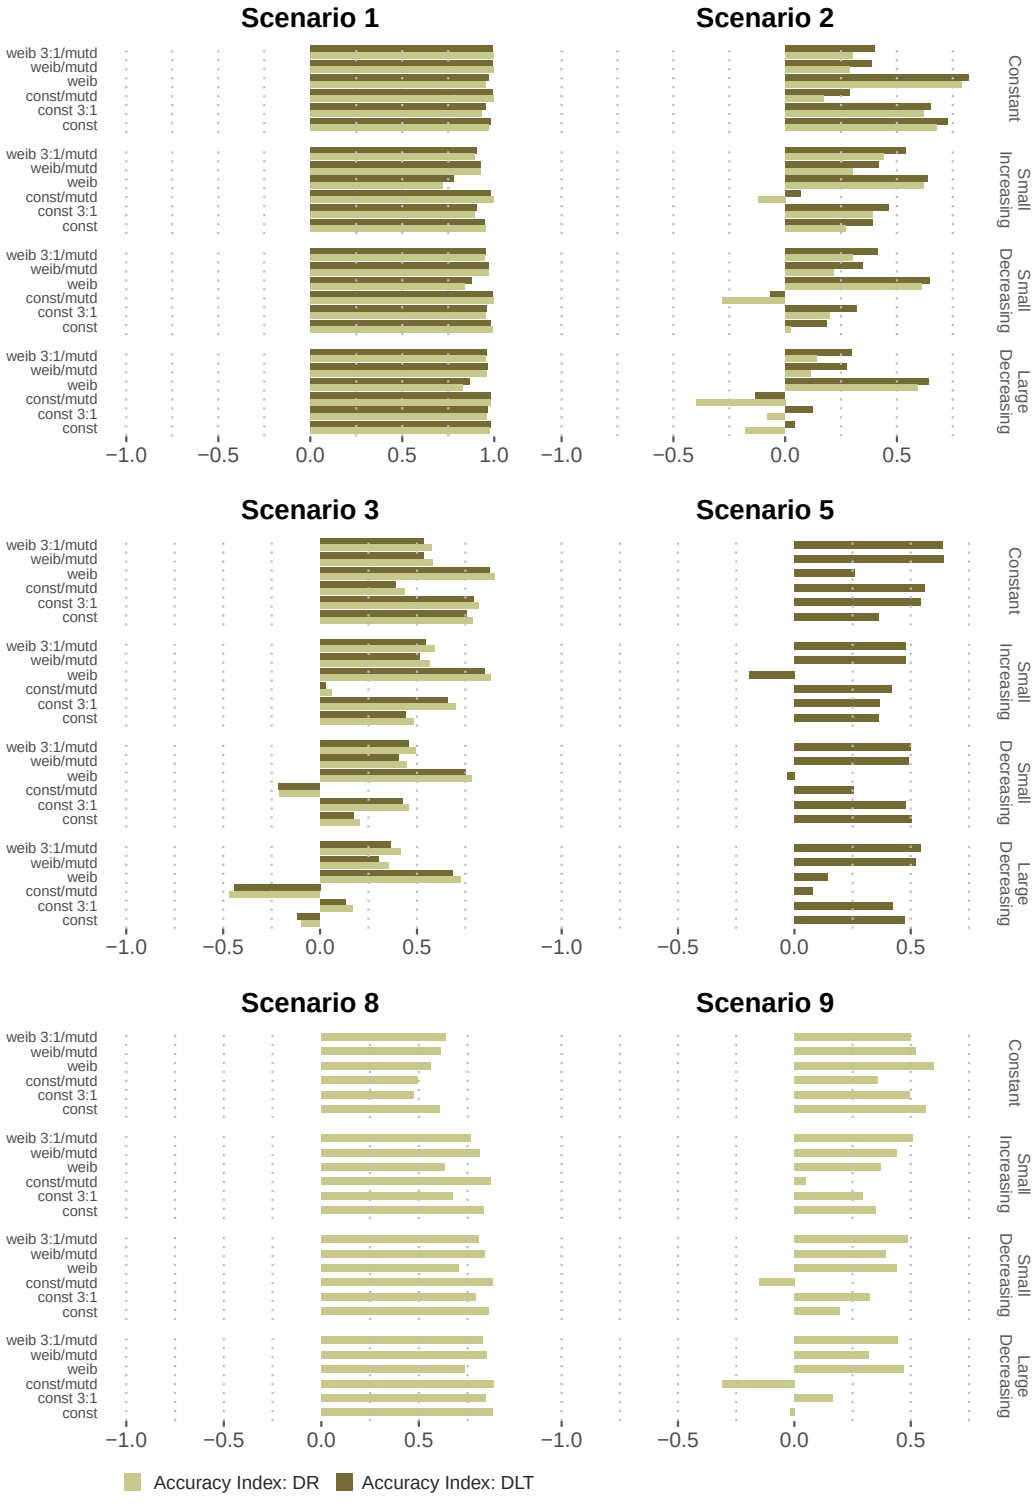

Figure A.8: Dose Reduction and DLT Accuracy Index values for each candidate model over all data generating hazards ($n = 60$).

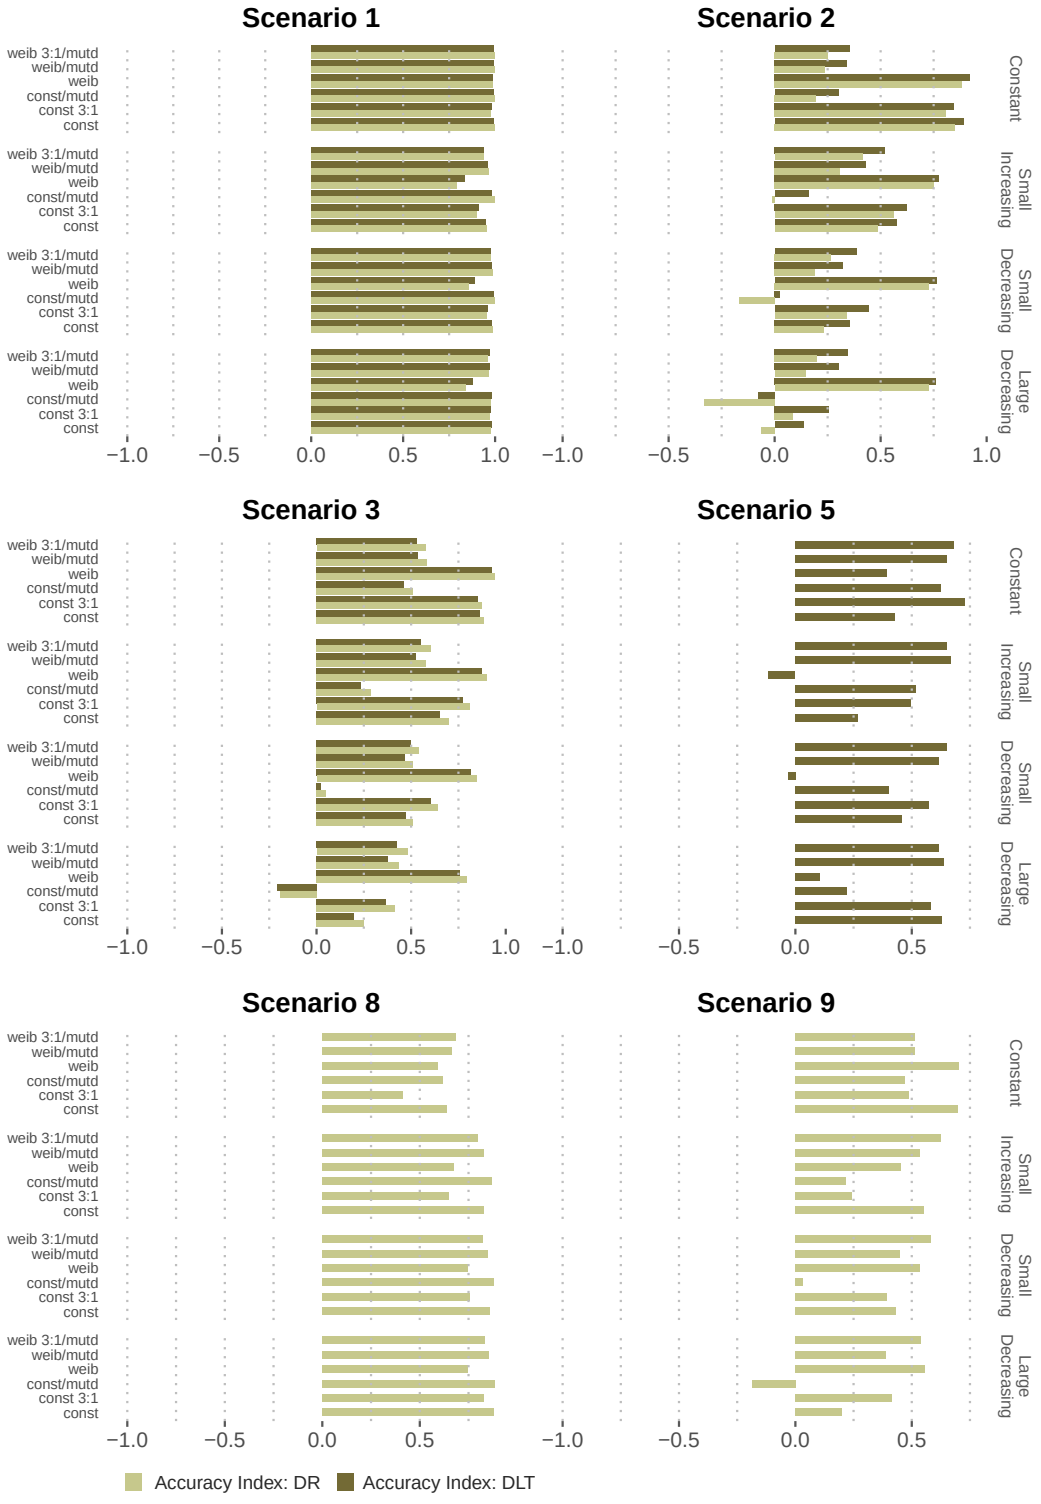

Figure A.9: Dose Reduction and DLT Accuracy Index values for each candidate model over all data generating hazards ($n = 100$).

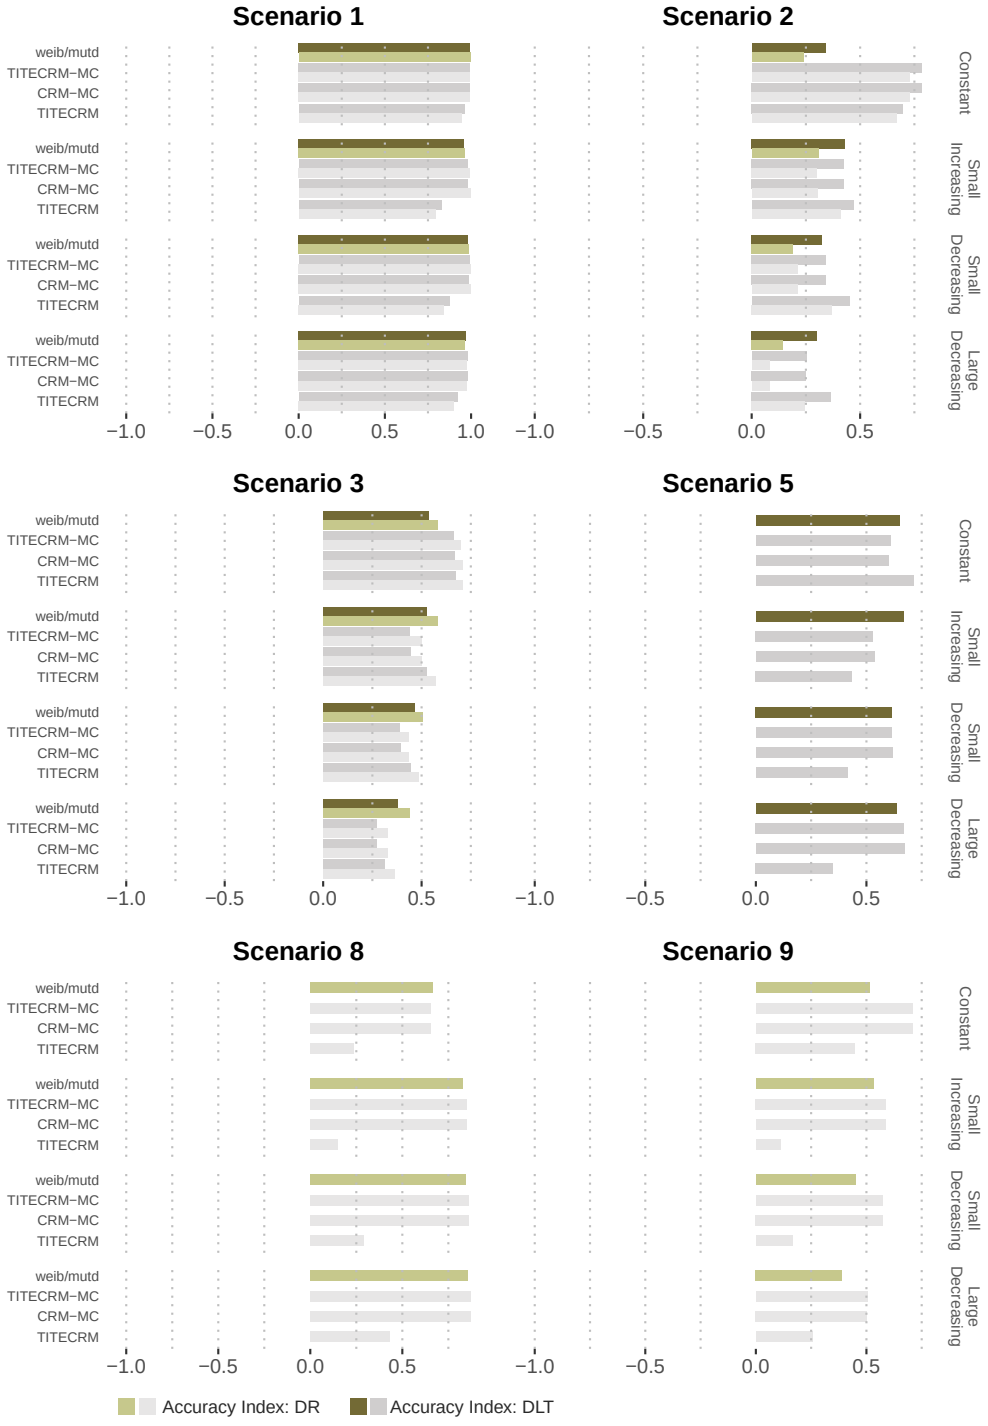

Figure A.10: Dose Reduction (light brown/light gray) and DLT (brown/gray) Accuracy Index values for Weibull-Skeleton + MUTD candidate model and comparator models ($n = 30$).

Hazards	Scenario	Targets	Parameters					True Probabilities									
								Dose Reduction					DLT				
			alpha	lambda1	lambda2	beta0	beta1	Dose1	Dose2	Dose3	Dose4	Dose5	Dose1	Dose2	Dose3	Dose4	Dose5
Constant	Scenario 7	0.5/0.3	1	0.248	0.0545		-0.5	65	79	82	82	82	30	48	66	81	89
Constant	Scenario 8	0.5/0.3	1	0.066	0.0177		-0.5	28	46	65	74	78	10	18	30	43	54
Constant	Scenario 9	0.5/0.3	1	0.027	0.0077		-0.5	13	24	40	54	65	4	8	15	22	30
Constant	Scenario 7	0.4/0.2	1	0.174	0.036		-0.5	55	74	82	83	83	20	34	51	67	77
Constant	Scenario 8	0.4/0.2	1	0.0447	0.011		-0.5	20	35	55	68	76	6	11	20	30	39
Constant	Scenario 9	0.4/0.2	1	0.0181	0.0047		-0.5	9	17	30	43	55	3	5	9	14	20
Constant	Scenario 4	0.5/0.3	1	0.1	0.061		-0.5	35	51	60	62	62	30	50	71	85	92
Constant	Scenario 5	0.5/0.3	1	0.0255	0.0175		-0.5	12	21	35	45	53	9	17	30	43	56
Constant	Scenario 6	0.5/0.3	1	0.0105	0.0073		-0.5	5	10	18	27	36	4	7	14	22	30
Constant	Scenario 4	0.4/0.2	1	0.062	0.04		-0.5	25	40	54	59	61	20	36	56	72	82
Constant	Scenario 5	0.4/0.2	1	0.0156	0.0108		-0.5	8	14	25	35	44	6	11	20	30	41
Constant	Scenario 6	0.4/0.2	1	0.0063	0.0045		-0.5	3	6	12	18	25	2	5	9	14	20
Constant	Scenario 1	0.5/0.3	1	0.162	0.058		-0.5	50	66	73	74	74	30	49	69	83	91
Constant	Scenario 2	0.5/0.3	1	0.042	0.0175		-0.5	19	33	50	61	67	9	17	30	43	54
Constant	Scenario 3	0.5/0.3	1	0.017	0.0074		-0.5	8	16	28	39	50	4	8	14	22	30
Constant	Scenario 1	0.4/0.2	1	0.111	0.038		-0.5	40	58	71	74	74	20	35	53	69	80
Constant	Scenario 2	0.4/0.2	1	0.029	0.011		-0.5	14	25	41	54	63	6	11	20	30	40
Constant	Scenario 3	0.4/0.2	1	0.0113	0.0045		-0.5	6	11	20	30	40	2	5	9	14	20
Small Inc.	Scenario 7	0.5/0.3	1.1	0.0156	0.0049	-0.5	2	65	69	74	76	76	30	34	40	49	60
Small Inc.	Scenario 8	0.5/0.3	1.1	0.0088	0.00295	-0.5	2	50	56	65	71	73	21	24	30	38	46
Small Inc.	Scenario 9	0.5/0.3	1.1	0.0047	0.0015	-0.5	2	34	39	50	60	65	12	14	19	25	30
Small Inc.	Scenario 7	0.4/0.2	1.1	0.0101	0.00284	-0.5	2	55	61	70	75	77	20	23	28	36	44
Small Inc.	Scenario 8	0.4/0.2	1.1	0.0056	0.0017	-0.5	2	38	44	55	65	69	13	16	20	27	32
Small Inc.	Scenario 9	0.4/0.2	1.1	0.00305	0.00087	-0.5	2	24	28	38	49	55	7	9	12	17	20
Small Inc.	Scenario 4	0.5/0.3	1.1	0.0058	0.0043	-0.5	2	35	40	47	53	55	30	35	43	53	62
Small Inc.	Scenario 5	0.5/0.3	1.1	0.00315	0.0024	-0.5	2	23	27	35	43	47	19	22	30	39	45
Small Inc.	Scenario 6	0.5/0.3	1.1	0.00173	0.0013	-0.5	2	14	17	23	31	35	11	13	19	25	30
Small Inc.	Scenario 4	0.4/0.2	1.1	0.00345	0.00255	-0.5	2	25	29	37	45	49	20	23	31	40	47
Small Inc.	Scenario 5	0.4/0.2	1.1	0.0019	0.0014	-0.5	2	15	18	25	33	37	12	14	20	27	32
Small Inc.	Scenario 6	0.4/0.2	1.1	0.00104	0.00077	-0.5	2	9	11	15	21	25	7	8	12	17	20
Small Inc.	Scenario 1	0.5/0.3	1.1	0.0097	0.0046	-0.5	2	50	55	62	66	67	31	34	42	51	60
Small Inc.	Scenario 2	0.5/0.3	1.1	0.0054	0.0026	-0.5	2	36	41	50	59	62	20	23	30	38	45
Small Inc.	Scenario 3	0.5/0.3	1.1	0.0029	0.00138	-0.5	2	23	26	35	45	50	12	14	19	25	30
Small Inc.	Scenario 1	0.4/0.2	1.1	0.0063	0.0027	-0.5	2	40	45	55	63	66	20	23	30	38	45
Small Inc.	Scenario 2	0.4/0.2	1.1	0.00345	0.00153	-0.5	2	26	30	40	50	55	13	15	20	27	32
Small Inc.	Scenario 3	0.4/0.2	1.1	0.0019	0.00082	-0.5	2	16	19	26	35	40	7	9	12	17	20

Figure A.11: Parameter settings for generating data from Constant and Small Increasing hazards. Small Increasing uses Cox-Weibull style hazards, while Constant uses the proposed model to give an example of assuming a "correct" model, which is why only one Beta parameter is given.

Hazards	Scenario	Targets	Parameters					True Probabilities									
								Dose Reduction					DLT				
			alpha	lambda1	lambda2	beta0	beta1	Dose1	Dose2	Dose3	Dose4	Dose5	Dose1	Dose2	Dose3	Dose4	Dose5
Large Dec.	Scenario 7	0.5/0.3	0.5	0.167	0.052	-0.5	2	65	69	74	76	76	30	33	39	49	59
Large Dec.	Scenario 8	0.5/0.3	0.5	0.093	0.031	-0.5	2	50	55	65	71	73	21	24	30	37	45
Large Dec.	Scenario 9	0.5/0.3	0.5	0.0504	0.0161	-0.5	2	34	39	50	60	65	12	14	19	25	30
Large Dec.	Scenario 7	0.4/0.2	0.5	0.1085	0.031	-0.5	2	55	60	70	75	77	20	23	29	36	44
Large Dec.	Scenario 8	0.4/0.2	0.5	0.06	0.0178	-0.5	2	38	44	55	65	69	13	15	20	26	32
Large Dec.	Scenario 9	0.4/0.2	0.5	0.0326	0.0093	-0.5	2	24	28	38	49	55	7	9	12	17	20
Large Dec.	Scenario 4	0.5/0.3	0.5	0.0613	0.0454	-0.5	2	35	39	47	53	55	30	34	43	53	61
Large Dec.	Scenario 5	0.5/0.3	0.5	0.034	0.026	-0.5	2	23	27	35	43	47	19	23	30	39	46
Large Dec.	Scenario 6	0.5/0.3	0.5	0.0185	0.0139	-0.5	2	14	17	23	31	35	11	13	19	25	30
Large Dec.	Scenario 4	0.4/0.2	0.5	0.0374	0.0275	-0.5	2	25	29	37	45	49	20	24	31	40	47
Large Dec.	Scenario 5	0.4/0.2	0.5	0.0204	0.0153	-0.5	2	15	18	25	33	37	12	14	20	27	32
Large Dec.	Scenario 6	0.4/0.2	0.5	0.0112	0.0082	-0.5	2	9	11	15	21	25	7	8	12	17	20
Large Dec.	Scenario 1	0.5/0.3	0.5	0.104	0.049	-0.5	2	50	55	62	66	67	30	34	41	51	60
Large Dec.	Scenario 2	0.5/0.3	0.5	0.057	0.028	-0.5	2	35	40	50	58	62	20	23	30	38	45
Large Dec.	Scenario 3	0.5/0.3	0.5	0.031	0.015	-0.5	2	22	26	35	45	50	12	14	19	25	30
Large Dec.	Scenario 1	0.4/0.2	0.5	0.068	0.029	-0.5	2	40	45	55	63	66	20	23	30	38	45
Large Dec.	Scenario 2	0.4/0.2	0.5	0.037	0.0165	-0.5	2	26	30	40	50	55	13	15	20	27	32
Large Dec.	Scenario 3	0.4/0.2	0.5	0.0202	0.0087	-0.5	2	16	19	26	35	40	7	9	12	17	20
Small Dec.	Scenario 7	0.5/0.3	0.8	0.0508	0.0159	-0.5	2	65	69	74	76	76	30	33	40	49	59
Small Dec.	Scenario 8	0.5/0.3	0.8	0.0288	0.0097	-0.5	2	50	56	65	71	73	21	24	30	38	46
Small Dec.	Scenario 9	0.5/0.3	0.8	0.0154	0.0049	-0.5	2	34	39	50	60	65	12	14	19	25	30
Small Dec.	Scenario 7	0.4/0.2	0.8	0.0332	0.0094	-0.5	2	55	61	70	75	77	20	23	28	36	44
Small Dec.	Scenario 8	0.4/0.2	0.8	0.0183	0.0054	-0.5	2	38	44	55	65	69	13	15	20	26	31
Small Dec.	Scenario 9	0.4/0.2	0.8	0.01	0.0028	-0.5	2	24	28	38	50	55	7	9	12	16	20
Small Dec.	Scenario 4	0.5/0.3	0.8	0.0188	0.014	-0.5	2	35	39	47	53	55	30	35	43	53	61
Small Dec.	Scenario 5	0.5/0.3	0.8	0.0103	0.0078	-0.5	2	23	27	35	43	47	19	22	30	39	45
Small Dec.	Scenario 6	0.5/0.3	0.8	0.0056	0.0042	-0.5	2	14	17	23	31	35	11	13	18	25	30
Small Dec.	Scenario 4	0.4/0.2	0.8	0.0114	0.0084	-0.5	2	25	29	37	45	49	20	24	31	40	47
Small Dec.	Scenario 5	0.4/0.2	0.8	0.0062	0.0046	-0.5	2	15	18	25	33	37	12	14	20	27	32
Small Dec.	Scenario 6	0.4/0.2	0.8	0.0034	0.0025	-0.5	2	9	11	15	21	25	7	8	12	17	20
Small Dec.	Scenario 1	0.5/0.3	0.8	0.0315	0.0148	-0.5	2	50	55	62	66	67	30	34	41	51	60
Small Dec.	Scenario 2	0.5/0.3	0.8	0.0175	0.0087	-0.5	2	35	40	50	58	61	20	23	30	39	46
Small Dec.	Scenario 3	0.5/0.3	0.8	0.0095	0.0045	-0.5	2	23	26	35	45	50	12	14	19	25	30
Small Dec.	Scenario 1	0.4/0.2	0.8	0.0206	0.0088	-0.5	2	40	45	55	63	66	20	23	30	38	45
Small Dec.	Scenario 2	0.4/0.2	0.8	0.0113	0.005	-0.5	2	26	30	40	50	55	13	15	20	27	32
Small Dec.	Scenario 3	0.4/0.2	0.8	0.0062	0.0027	-0.5	2	16	19	26	35	40	7	9	12	17	20

Figure A.12: Parameter settings for generating data from Small Decreasing and Large Decreasing hazards. Both use Cox-Weibull style hazards.

Appendix B

Derivations and Specifications

B.1 Three-transition: transition probabilities

Constant-Skeleton model

$$p_{DR,j} = \frac{\lambda_1}{-\lambda_1 - \lambda_2} \times \exp\left((-\lambda_1 x_j - \lambda_2 x_j) * t_{max} - 1\right) \tag{B.1}$$

$$\begin{aligned} p_{DLT,j} = {} & \frac{\lambda_1}{-\lambda_1 - \lambda_2} \times \exp\left((-\lambda_1 x_j - \lambda_2 x_j) * t_{max} - 1\right) - \\ & \exp(-\lambda_2 x_{j-1}^{\exp\beta_1} t_{max}) \times \frac{\lambda_1 x_j}{-\lambda_1 x_j - \lambda_2 x_j + \lambda_2 x_{j-1}^{\exp(\beta_1)}} \times \\ & \exp\left((-\lambda_1 x_j - \lambda_2 x_j + \lambda_2 x_{j-1}^{\exp\beta_1}) * t_{max} - 1\right) + \\ & \frac{\lambda_2}{-\lambda_1 - \lambda_2} \times \exp\left((-\lambda_1 x_j - \lambda_2 x_j) * t_{max} - 1\right) \end{aligned} \tag{B.2}$$

Weibull-Skeleton model

$$p_{DR,j} = \frac{\lambda_1}{-\lambda_1 - \lambda_2} \times \exp\left((-\lambda_1 x_j - \lambda_2 x_j) * t_{max}^{\alpha} - 1\right) \tag{B.3}$$

$$\begin{aligned} p_{DLT,j} = {} & \frac{\lambda_1}{-\lambda_1 - \lambda_2} \times \exp\left((-\lambda_1 x_j - \lambda_2 x_j) * t_{max}^{\alpha} - 1\right) - \\ & \exp(-\lambda_2 x_{j-1}^{\exp\beta_1} t_{max}^{\alpha}) \times \frac{\lambda_1 x_j}{-\lambda_1 x_j - \lambda_2 x_j + \lambda_2 x_{j-1}^{\exp(\beta_1)}} \times \\ & \exp\left((-\lambda_1 x_j - \lambda_2 x_j + \lambda_2 x_{j-1}^{\exp\beta_1}) t_{max}^{\alpha} - 1\right) + \\ & \frac{\lambda_2}{-\lambda_1 - \lambda_2} \times \exp\left((-\lambda_1 x_j - \lambda_2 x_j) * t_{max}^{\alpha} - 1\right) \end{aligned} \tag{B.4}$$

B.2 Five-transition: hazards

The five-transition model needs five associated hazard functions to be specified. The following specifies each of the five uniquely with a Constant-Skeleton specification. As stated in the text, a frailty parameter γ_i could be multiplied on each hazard to induce dependence between observed events on the same patient.

$$h_1(t_1|x_j) = \lambda_1 x_{ij} \tag{B.5}$$

$$h_2(t_2|x_j) = \lambda_2 x_{ij} \tag{B.6}$$

$$h_3(t_2|t_1, x_j) = \lambda_2 x_{i,j-1}^{exp(\beta_1)} \tag{B.7}$$

$$h_4(t_2|t_1, x_j) = \lambda_1 x_{i,j-1}^{exp(\beta_1+\beta_2)} \tag{B.8}$$

$$h_5(t_2|t_1, x_j) = \lambda_2 x_{i,j-2}^{exp(\beta_1+\beta_2)} \tag{B.9}$$

The above model contains four paramters. An analogous Weibull-Skeleton specification would contain five parameters. If the setting dictates a five-transition model, but there is not much difference expected between the hazard of first Dose Reduction versus hazard of second Dose Reduction given first Dose Reduction, or hazard of DLT given one Dose Reduction versus hazard of DLT given two Dose Reductions, the previous hazards can be simplified. In this case, set $h_4 = h_1$ and/or $h_5 = h_3$ (but still changing doses). Thus, the five-transition setting is respected, but there are the same number of parameters to estimate as in the three-transition model.

B.3 Five-transition: likelihood and transition probability derivations

The observed data likelihood can again be separated into pieces. Let T_3 be the time to second Dose Reduction. likelihood

$$P(T_1 = t_1, T_2 = t_2, T_3 = t_3) \tag{B.10}$$

$$= P(t_3 < T_1 = t_1, T_2 = t_2, T_3 = t_3) + P(t_3 > T_1 = t_1, T_2 = t_2, T_3 = t_3) \tag{B.11}$$

$$= 0 + P(T_1 = t_1, T_2 = t_2, T_3 = t_3, T_1 < t_3) \tag{B.12}$$

$$= P(T_1 = t_1, T_2 = t_2, T_3 = t_3 | T_1 < t_3) P(T_1 < t_3) \tag{B.13}$$

$$= P(T_1 = t_1, T_2 = t_2, T_3 = t_3 | T_1 < t_3) \times 1 \tag{B.14}$$

Then,

$$P(T_1 = t_1, T_2 = t_2, T_3 = t_3 | T_1 < t_3) \tag{B.15}$$

$$= P(T_1 = t_1, T_2 = t_2, T_3 = t_3, T_1 > t_2, T_3 > t_2 | T_1 < t_3) \tag{B.16}$$

$$+ P(T_1 = t_1, T_2 = t_2, T_3 = t_3, T_2 > t_1, T_3 > t_2 | T_1 < t_3) \tag{B.17}$$

$$+ P(T_1 = t_1, T_2 = t_2, T_3 = t_3, T_2 > t_1, T_2 > t_3 | T_1 < t_3) \tag{B.18}$$

Let these three linear terms be defined as . These three linear terms or likelihood elements can be defined as f_c, f_d, and f_e and derived separately. Then, the elements are the foundation to generate the six likelihood paths in the five-transition model (as compared to four paths in the three-transition model).

Joint Density Piece 1. $f_c(t_1, t_2, t_3) = h_2(t_2)S_1(t_2)S_2(t_2)$

Derivation:

$$P(T_1 = t_1, T_2 = t_2, T_3 = t_3, T_1 > t_2, T_3 > t_2 | T_1 < t_3) \tag{B.19}$$

$$= P(T_1 = \infty, T_2 = t_2, T_3 = \infty, T_1 > t_2, T_3 > t_2 | T_1 < t_3) \tag{B.20}$$

$$= P(T_3 = \infty | T_1 = \infty, T_2 = t_2, T_1 > t_2, T_3 > t_2, T_1 < t_3) \times \tag{B.21}$$
$$\quad P(T_1 = \infty, T_2 = t_2, T_1 > t_2, T_3 > t_2 | T_1 < t_3)$$

$$= P(T_3 = \infty | T_1 = \infty, T_2 = t_2, T_1 > t_2, T_3 > t_2, T_1 < t_3) \times \tag{B.22}$$
$$\quad P(T_1 = \infty | T_2 = t_2, T_1 > t_2, T_3 > t_2, T_1 < t_3) \times$$
$$\quad P(T_2 = t_2, T_1 > t_2, T_3 > t_2, T_1 < t_3)$$

$$= 1 \times 1 \times \frac{P(T_1 > t_2, T_3 > t_2, T_2 = t_2 | T_1 < t_3)}{P(T_1 > t_2, T_3 > t_2, T_2 > t_2 | T_1 < t_3)} \times \tag{B.23}$$
$$\quad P(T_1 > t_2, T_3 > t_2, T_2 > t_2 | T_1 < t_3)$$

$$= \frac{P(T_2 = t_2 | T_1 > t_2, T_2 > t_2, T_1 < t_3)}{P(T_2 > t_2 | T_1 > t_2, T_2 > t_2, T_1 < t_3)} P(T_3 > t_2 | T_1 > t_2, T_2 > t_2, T_1 < t_3) \times \tag{B.24}$$
$$\quad P(T_1 > t_2, T_2 > t_2 | T_1 < t_3)$$

$$= h_2(t_2)S_1(t_2)S_2(t_2) \tag{B.25}$$

Joint Density Piece 2. $f_d(t_1,t_2,t_3) = h_1(t_1)S_1(t_1)S_2(t_1)h_3(t_2|t_1)\dfrac{S_3(t_2|t_1)}{S_3(t_1|t_1)}\dfrac{S_4(t_2|t_1)}{S_4(t_1|t_1)}$

Derivation:

$$P(T_1 = t_1, T_2 = t_2, T_3 = t_3, T_2 > t_1, T_3 > t_2 | T_1 < t_3)$$

$$P(T_1 = t_1, T_2 = t_2, T_3 = \infty, T_2 > t_1, T_3 > t_2 | T_1 < t_3) \tag{B.26}$$

$$= P(T_2 = t_2 | T_1 = t_1, T_3 = \infty, T_2 > t_1, T_3 > t_2, T_1 < t_3) \times \tag{B.27}$$

$$P(T_1 = t_1, T_3 = \infty, T_2 > t_1, T_3 > t_2 | T_1 < t_3)$$

$$= P(T_2 = t_2 | T_1 = t_1, T_3 = \infty, T_2 > t_1, T_3 > t_2, T_1 < t_3) \times \tag{B.28}$$

$$\frac{P(T_2 > t_2 | T_1 = t_1, T_3 = \infty, T_2 > t_1, T_3 > t_2, T_1 < t_3)}{P(T_2 > t_2 | T_1 = t_1, T_3 = \infty, T_2 > t_1, T_3 > t_2, T_1 < t_3)} \times$$

$$P(T_3 > \infty | T_1 = t_1, T_2 > t_1, T_3 > t_2, T_1 < t_3) \times$$

$$P(T_1 = t_1, T_2 > t_1, T_3 > t_2 | T_1 < t_3)$$

$$= h_3(t_2|t_1)\frac{S_3(t_2|t_1)}{S_3(t_1|t_1)} \times 1 \times P(T_1 = t_1, T_2 > t_1, T_3 > t_2 | T_1 < t_3) \tag{B.29}$$

$$= h_3(t_2|t_1)\frac{S_3(t_2|t_1)}{S_3(t_1|t_1)} P(T_3 > t_2 | T_1 = t_1, T_2 > t_1, T_1 < t_3) \times \tag{B.30}$$

$$P(T_1 = t_1, T_2 > t_1 | T_1 < t_3)$$

$$= h_3(t_2|t_1)\frac{S_3(t_2|t_1)}{S_3(t_1|t_1)}\frac{S_4(t_2|t_1)}{S_4(t_1|t_1)} P(T_1 = t_1, T_2 > t_1 | T_1 < t_3) \times \tag{B.31}$$

$$\frac{P(T_1 > t_1, T_2 > t_1 | T_1 < t_3)}{P(T_1 > t_1, T_2 > t_1 | T_1 < t_3)}$$

$$= h_1(t_1)S_1(t_1)S_2(t_1)h_3(t_2|t_1)\frac{S_3(t_2|t_1)}{S_3(t_1|t_1)}\frac{S_4(t_2|t_1)}{S_4(t_1|t_1)} \tag{B.32}$$

Joint Density Piece 3. $f_e(t_1,t_2,t_3) = h_1(t_1)S_1(t_1)S_2(t_1)h_4(t_3|t_1)\frac{S_3(t_3|t_1)}{S_3(t_1|t_1)}\frac{S_4(t_3|t_1)}{S_4(t_1|t_1)}\times$
$h_5(t_2|t_3,t_1)\frac{S_5(t_2|t_3,t_1)}{S_5(t_3|t_1,t_3)}$

Derivation:

$P(T_1 = t_1, T_2 = t_2, T_3 = t_3, T_2 > t_1, T_2 > t_3 | T_1 < t_3)$

$= P(T_2 = t_2|T_1 = t_1, T_3 = t_3, T_2 > t_1, T_2 > t_3, T_1 < t_3)\times$ (B.33)
$\frac{P(T_2 > t_2|T_1 = t_1, T_3 = t_3, T_2 > t_3, T_2 > t_1, T_1 < t_3)}{P(T_2 > t_2|T_1 = t_1, T_3 = t_3, T_2 > t_3, T_2 > t_1, T_1 < t_3)}\times$
$\frac{P(T_3 > t_3|T_1 = t_1, T_2 > t_3, T_2 > t_1, T_1 < t_3)}{P(T_3 > t_3|T_1 = t_1, T_2 > t_3, T_2 > t_1, T_1 < t_3)}\times$
$P(T_3 = t_3|T_1 = t_1, T_2 > t_1, T_2 > t_3, T_1 < t_3)\times$
$P(T_1 = t_1, T_2 > t_1, T_2 > t_3 | T_1 < t_3)$

$= h_5(t_2|t_1,t_3)\frac{S_5(t_2|t1,t_3}{S_5(t_3|t_1,t_3}h_4(t_3|t_1,t_3)\frac{S_4(t_3|t_1,t_3)}{S_4(t_1|t_1,t_3)}\times$ (B.34)
$P(T_2 > t_3|T_1 = t_1, T_2 > t_1, T_1 < t_3)P(T_1 = t_1, T_2 > t_1|T_1 < t_3)$

$= h_5(t_2|t_1,t_3)\frac{S_5(t_2|t_1,t_3)}{S_5(t_3|t_1,t_3)}h_4(t_3|t_1,t_3)\frac{S_4(t_3|t_1,t_3)}{S_4(t_1|t_1,t_3)}\frac{S_3(t_3|t_1,t_3)}{S_3(t_1|t_1,t_3)}\times$ (B.35)
$\frac{P(T_1 > t_1, T_2 > t_1|T_1 < t_3)}{P(T_1 > t_1, T_2 > t_1|T_1 < t_3)}P(T_1 = t_1, T_2 > t_1|T_1 < t_3)$

$= h_1(t_1)S_1(t_1)S_2(t_1)h_4(t_3|t_1)\frac{S_3(t_3|t_1)}{S_3(t_1|t_1)}\frac{S_4(t_3|t_1)}{S_4(t_1|t_1)}h_5(t_2|t_3,t_1)\frac{S_5(t_2|t_3,t_1)}{S_5(t_3|t_1,t_3)}$ (B.36)

1. Censored before any toxicity:

$L_1 = P(T_1 > C, T_2 > C, T_3 > C | T_1 < t_3) = S_1(C)S_2(C)$

Derivation:

$$P(T_1 > C, T_2 > C, T_3 > C | T_1 < t_3) \tag{B.37}$$

$$= P(T_1 > C, T_2 > C | T_1 < t_3) \tag{B.38}$$

$$= S_1(C)S_2(C)$$

Again, direct result from joint survivor function as explained in (3.14).

2. DLT before Dose Reduction:

$L_2 = P(T_1 = \infty, T_3 = \infty, T_2 = t_2, T_1 > t_2, T_3 > t_2 | T_1 < t_3) = h_2(t_2)S_1(t_2)S_2(t_2)$

Derivation:

$$P(T_1 = \infty, T_3 = \infty, T_2 = t_2, T_1 > t_2, T_3 > t_2 | T_1 < t_3) \tag{B.39}$$

$$= P(T_1 = t_1, T_2 = t_2, T_3 = t_3, T_1 > t_2, T_3 > t_2 | T_1 < t_3) \tag{B.40}$$

$$= f_c(t_1, t_2, t_3) \tag{B.41}$$

$$= h_2(t_2)S_1(t_2)S_2(t_2) \tag{B.42}$$

3. Censored after 1 Dose Reduction:

$$L_3 = P(T_1 = t_1, T_3 > C, T_2 > C, T_2 > t_3, T_2 > t_1 | T_1 < t_3)$$
$$= h_1(t_1)S_1(t_1)S_2(t_1)\frac{S_4(t_3|t_1)}{S_4(t_1|t_1)}\frac{S_3(t_3|t_1)}{S_3(t_1|t_1)}$$

Derivation:

$$P(T_1 = t_1, T_3 > C, T_2 > C, T_2 > t_3, T_2 > t_1 | T_1 < t_3)$$
$$= \int_C^\infty f_d(t_1, s, t_3)ds \tag{B.43}$$

(except h_4 is also included in f_d as it is a possible event in this scenario, (B.44)

that previously f_d specifying a DLT did not contain)

$$= \frac{h_1(t_1)S_1(t_1)S_2(t_1)}{S_4(t_1|t_1)S_3(t_1|t_1)}\int_C^\infty h_3(s|t_1)h_4(s|t_1)\exp(-H_3(s|t_1))\exp(-H_4(s|t_1))ds \tag{B.45}$$
$$= h_1(t_1)S_1(t_1)S_2(t_1)\frac{S_4(C|t_1)}{S_4(t_1|t_1)}\frac{S_3(C|t_1)}{S_3(t_1|t_1)} \tag{B.46}$$

4. DLT after 1 Dose Reduction:

$$L_4 = P(T_1 = t_1, T_3 = \infty, T_2 = t_2, T_2 > t_1, T_3 > t_2 | T_1 < t_3)$$
$$= h_1(t_1)S_1(t_1)S_2(t_1)h_3(t_2|t_1)\frac{S_3(t_2|t_1)}{S_3(t_1|t_1)}\frac{S_4(t_2|t_1)}{S_4(t_1|t_1)}$$

Derivation:

$$P(T_1 = t_1, T_3 = \infty, T_2 = t_2, T_2 > t_1, T_3 > t_2 | T_1 < t_3)$$
$$= P(T_1 = t_1, T_3 = t_3, T_2 = t_2, T_2 > t_1, T_3 > t_2 | T_1 < t_3)$$
$$= f_d(t_1, t_2, t_3) \tag{B.47}$$
$$= h_1(t_1)S_1(t_1)S_2(t_1)h_3(t_2|t_1)\frac{S_3(t_2|t_1)}{S_3(t_1|t_1)}\frac{S_4(t_2|t_1)}{S_4(t_1|t_1)} \tag{B.48}$$

5. Censored after 2 Dose Reductions:

$$L_6 = P(T_1 = t_1, T_3 = t_3, T_2 > C_i, T_2 > t_3, T_2 > t_1 | T_1 < t_3)$$
$$= h_1(t_1)S_1(t_1)S_2(t_1)\frac{S_4(t_3|t_1)}{S_4(t_1|t_1)}\frac{S_3(t_3|t_1)}{S_3(t_1|t_1)}h_4(t_3|t_1)\frac{S_5(C_i|t_1,t_3)}{S_5(t_3|t_1,t_3)}$$

Derivation:

$$P(T_1 = t_1, T_3 = t_3, T_2 > C_i, T_2 > t_3, T_2 > t_1 | T_1 < t_3)$$
$$= \int_C^\infty f_e(t_1, t_3, s)ds \tag{B.49}$$
$$= h_1(t_1)S_1(t_1)S_2(t_1)\frac{S_4(t_3|t_1)}{S_4(t_1|t_1)}\frac{S_3(t_3|t_1)}{S_3(t_1|t_1)}h_4(t_3|t_1)\frac{1}{S_5(t_3|t_1)}\int_C^\infty h_5(s|t_1,t_3)\exp(-H_5(s|t_1,t_3))ds \tag{B.50}$$
$$= h_1(t_1)S_1(t_1)S_2(t_1)\frac{S_4(t_3|t_1)}{S_4(t_1|t_1)}\frac{S_3(t_3|t_1)}{S_3(t_1|t_1)}h_4(t_3|t_1)\frac{S_5(C|t_1,t_3)}{S_5(t_3|t_1,t_3)} \tag{B.51}$$

6. DLT after 2 Dose Reductions:

$$L_6 = P(T_1 = t_1, T_3 = t_3, T_2 = t_2, T_2 > t_3, T_2 > t_1 | T_1 < t_3)$$
$$= h_1(t_1)S_1(t_1)S_2(t_1)h_4(t_3|t_1)\frac{S_3(t_3|t_1)}{S_3(t_1|t_1)}\frac{S_4(t_3|t_1)}{S_4(t_1|t_1)}h_5(t_2|t_3,t_1)\frac{S_5(t_2|t_3,t_1)}{S_5(t_3|t_1,t_3)}$$

Derivation:

$$P(T_1 = t_1, T_3 = t_3, T_2 = t_{2i}, T_2 > t_{1i}, T_2 > t_3 | T_1 < t_3)$$
$$= f_e(t_1, t_2, t_3) \tag{B.52}$$
$$= h_1(t_1)S_1(t_1)S_2(t_1)h_4(t_3|t_1)\frac{S_3(t_3|t_1)}{S_3(t_1|t_1)}\frac{S_4(t_3|t_1)}{S_4(t_1|t_1)}h_5(t_2|t_3,t_1)\frac{S_5(t_2|t_3,t_1)}{S_5(t_3|t_1,t_3)} \tag{B.53}$$

The likelihood could also be derived using the Counting Process method. The result will be the same.

Transition Probabilities

In a five-transition model, the two most important transition probabilities are still Healthy → DR ($p_{DR,j}$) and Healthy → DLT via any path ($p_{DLT,j}$). Healthy → DR has the same specification as the three-transition model.

Healthy → DR: $P_{t_c}(T_1 < t, T_2 > t, T_3 > t | T_1 > t_c, T_2 > t_c, T_1 < t_3)$ (B.54)

$$= \frac{\int_{t_c}^{t} h_1(u) S_1(u) S_2(u)}{S_1(t_c) S_2(t_c)}$$

Derivation:

$P_{t_c}(T_1 < t, T_2 > t, T_3 > t | T_1 > t_c, T_2 > t_c, T_1 < t_3)$

$$= \int_{t_c}^{t} P(T_1 = u, T_2 > u, T_3 > u | T_1 > t_c, T_2 > t_c, T_1 < t_3) du \tag{B.55}$$

$$= \frac{\int_{t_c}^{t} P(T_1 = u, T_2 > u, T_3 > u | T_1 < t_3) du}{P(T_1 > t_c, T_2 > t_c | T_1 < t_3)} \tag{B.56}$$

$$= \frac{\int_{t_c}^{t} P(T_1 = u, T_2 > u, T_3 > u | T_1 < t_3) du}{S_t(t_c) S_2(t_c)} \tag{B.57}$$

$$= \frac{\int_{t_c}^{t} P(T_1 = u, T_2 > u, T_3 > u | T_1 < t_3) \dfrac{P(T_1 > u, T_2 > u, T_3 > u | T_1 < t_3)}{P(T_1 > u, T_2 > u, T_3 > u | T_1 < t_3)} du}{S_t(t_c) S_2(t_c)} \tag{B.58}$$

$$= \frac{\int_{t_c}^{t} h_1(u) S_1(u) S_2(u)}{S_1(t_c) S_2(t_c)} \tag{B.59}$$

Healthy → DLT via any path: $Pt_c(T_2 < t | T_1 < t, T_3 < t, T_1 < t_3, T_1 > t_c, T_2 > t_c)$

$$P(\text{Healthy} \rightarrow \text{DLT directly}) + P(\text{Healthy} \rightarrow \text{DR1} \rightarrow \text{DLT}) + \tag{B.60}$$

$$P(\text{Healthy} \rightarrow \text{DR1} \rightarrow \text{DR2} \rightarrow \text{DLT}) = a + b + c$$

Derivation:

$$a = P(T_2 < t, T_1 > t, T_3 > t | T_1 > t_c, T_2 > t_c, T_1 < t_3) \tag{B.61}$$

$$= \int_{t_c}^{t} P(T_2 = u, T_1 > u, T_3 > u | T_1 > t_c, T_2 > t_c, T_1 < t_3) du \tag{B.62}$$

$$= \int_{t_c}^{t} P(T_2 = u, T_1 = \infty, T_3 = \infty, T_1 > u, T_3 > u | T_1 > t_c, T_2 > t_c, T_1 < t_3) du \tag{B.63}$$

$$= \frac{\int_{t_c}^{t} P(T_2 = u, T_1 = \infty, T_3 = \infty, T_1 > u, T_3 > u, T_1 > t_c, T_2 > t_c | T_1 < t_3) du}{P(T_1 > t_c, T_2 > t_c | T_1 < t_3)} \tag{B.64}$$

$$= \frac{\int_{t_c}^{t} P(T_2 = u, T_1 = \infty, T_3 = \infty, T_1 > u, T_3 > u | T_1 < t_3) du}{P(T_1 > t_c, T_2 > t_c | T_1 < t_3)} \tag{B.65}$$

From joint density piece 1,

$$= \frac{\int_{t_c}^{t} h_2(u) S_1(u) S_2(u) du}{S_1(t_c) S_2(t_c)} \tag{B.66}$$

$$b = P(\text{Healthy} \rightarrow \text{DR1})P(\text{DR1} \rightarrow \text{DLT} \mid \text{Healthy} \rightarrow \text{DR1}) \tag{B.67}$$

$$= P(\text{Healthy} \rightarrow \text{DR1})P(T_2 < t, T_3 > t | T_1 < u, T_2 > u, T_3 > u, T_1 < t_3) \tag{B.68}$$

$$= P(\text{Healthy} \rightarrow \text{DR1}) \int_u^t P(T_2 = r, T_3 > r | T_1 < u, T_2 > u, T_3 > u, T_1 < t_3) dr \tag{B.69}$$

$$= P(\text{Healthy} \rightarrow \text{DR1}) \frac{\int_u^t P(T_2 = r, T_3 > r, T_2 > u, T_3 > u | T_1 < u, T_1 < t_3) dr}{P(T_2 > u, T_3 > u | T_1 < u, T_1 < t_3)} \tag{B.70}$$

$$= P(\text{Healthy} \rightarrow \text{DR1}) \frac{\int_u^t P(T_2 = r, T_3 > r | T_1 < u, T_1 < t_3) dr}{S_3(u) S_4(u)} \tag{B.71}$$

$$= P(\text{Healthy} \rightarrow \text{DR1}) \times \tag{B.72}$$

$$\frac{\int_u^t P(T_2 = r, T_3 > r | T_1 < u, T_1 < t_3) \dfrac{P(T_2 > r, T_3 > r | T_1 < u, T_1 < t_3)}{P(T_2 > r, T_3 > r | T_1 < u, T_1 < t_3)} dr}{S_3(u) S_4(u)}$$

$$= \frac{\int_{t_c}^t h_1(u) S_1(u) S_2(u) \dfrac{\int_u^t h_3(r) S_4(r) S_3(r) dr}{S_3(u) S_4(u)} du}{S_1(t_c) S_2(t_c)} \tag{B.73}$$

$$c = P(\text{Healthy} \rightarrow \text{DR1})P(\text{DR1} \rightarrow \text{DR2} \mid \text{Healthy} \rightarrow \text{DR1}) \times \tag{B.74}$$

$$P(\text{DR2} \rightarrow \text{DLT} \mid \text{Healthy} \rightarrow \text{DR1}, \text{DR1} \rightarrow \text{DLT})$$

$P(\text{DR1} \rightarrow \text{DR2} \mid \text{Healthy} \rightarrow \text{DR1})$ follows similarly from $P(\text{DR1} \rightarrow \text{DLT} \mid \text{Healthy} \rightarrow \text{DR1})$ previously. $P(\text{DR2} \rightarrow \text{DLT} \mid \text{Healthy} \rightarrow \text{DR1}, \text{DR1} \rightarrow \text{DLT})$ follows similarly from Putter, Fiocco, and Geskus 2007, based in the fact there is only one state in

which to transition.

$$= P(\text{Healthy} \rightarrow \text{DR1})P(\text{DR1} \rightarrow \text{DR2} \mid \text{Healthy} \rightarrow \text{DR1})(1 - \frac{S_5(t)}{S_5(r)}) \tag{B.75}$$

$$= \frac{\int_{t_c}^{t} h_1(u)S_1(u)S_2(u)\frac{\int_u^t h_3(r)S_4(r)S_3(r)(1 - \frac{S_5(t)}{S_5(r)})dr}{S_3(u)S_4(u)}du}{S_1(t_c)S_2(t_c)} \tag{B.76}$$

Thus, the full transition probability $a + b + c$,

$$P_{t_c}(T_2 < t | T_1 < t, T_3 < t, T_1 < t_3, T_1 > t_c, T_2 > t_c)$$

$$= a + b + c$$

$$= \frac{\int_{t_c}^{t} h_2(u)S_1(u)S_2(u)du}{S_1(t_c)S_2(t_c)} + \tag{B.77}$$

$$\frac{\int_{t_c}^{t} h_1(u)S_1(u)S_2(u)\frac{\int_u^t h_3(r)S_4(r)S_3(r)dr}{S_3(u)S_4(u)}du}{S_1(t_c)S_2(t_c)} +$$

$$\frac{\int_{t_c}^{t} h_1(u)S_1(u)S_2(u)\frac{\int_u^t h_3(r)S_4(r)S_3(r)(1 - \frac{S_5(t)}{S_5(r)})dr}{S_3(u)S_4(u)}du}{S_1(t_c)S_2(t_c)}$$