



HAL
open science

Méthodes statistiques pour les essais de phase I/II de thérapies moléculaires ciblées en cancérologie

Maria Athina Altzerinakou

► **To cite this version:**

Maria Athina Altzerinakou. Méthodes statistiques pour les essais de phase I/II de thérapies moléculaires ciblées en cancérologie. Cancer. Université Paris Saclay (COmUE), 2018. Français. NNT : 2018SACLS375 . tel-02303040

HAL Id: tel-02303040

<https://theses.hal.science/tel-02303040>

Submitted on 2 Oct 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Statistical methods for phase I/II trials of molecularly targeted agents in oncology

Thèse de doctorat de l'Université Paris-Saclay
préparée à l'Université Paris-Sud

Ecole doctorale n°570 Santé Publique (EDSP)
Spécialité de doctorat : Santé publique - biostatistiques

Thèse présentée et soutenue à Villejuif, le Vendredi 12 Octobre 2018, par

MARIA-ATHINA ALTZERINAKOU

Composition du Jury :

Aurélien Latouche Professeur, Conservatoire national des arts et métiers (CNAM)	Président
Helena Geys Professeur, University of Hasselt	Rapporteur
Cécile Proust-Lima Chargée de Recherche, Institut de santé publique, d'épidémiologie et de développement (ISPED)	Rapporteur
David Perol Docteur en Médecine, Centre de lutte contre le cancer Léon Bérard	Examineur
Xavier Paoletti Docteur, Gustave Roussy Service de biostatistique et d'épidémiologie; INSERM CESP OncoStat; Université Paris-Saclay	Directeur de thèse

In memory of my father Panagiotis Altzerinakos

Acknowledgments

Je remercie Pr. Aurélien Latouche d'avoir accepté de présider mon jury de thèse. Je remercie également grandement Pr. Helena Geys et Pr. Cécile Proust-Lima d'avoir accepté d'être rapporteurs de ma thèse. Je les remercie pour tous leurs précieux commentaires et suggestions. Enfin, ma gratitude s'adresse également à Dr. David Perol d'avoir accepté d'examiner ce travail. Encore merci aux membres de mon jury d'avoir fait le déplacement pour ma soutenance.

Je tiens à remercier Xavier Paoletti d'avoir dirigé ce travail pendant trois années. Je pense avoir eu beaucoup de chance d'être encadrée par un directeur de thèse de cette qualité tant sur le plan professionnel que personnel. Thank you for the guidance, the teaching and the many many hours you devoted to and for me. And of course thank you for giving me the probably once in a life time opportunity to join a program like IDEAS. Ce fut un plaisir de travailler avec toi et je t'en remercie grandement!

I would like to thank the IDEAS training network and through that the European Union that funded this project.

Bien évidemment, je remercie l'ensemble du service de biostatistique et d'épidémiologie de Gustave Roussy pour tout ce qu'ils ont fait pour moi. Je remercie Ellen Benhamou et Stefan Michiels de m'avoir accueillie au sein du service et de l'équipe OncoStat.

I would also like to thank Richardus Vonk for hosting me in Bayer and for doing everything possible to make me feel welcome.

A big thank you to my Gustave Roussy colleagues. I will start from Federico and Béranger who first welcomed me and showed me how nice and cool it is to work with this team. Even bigger thank you to Federico for answering the millions of questions I had during these three years and for all the guidance he provided me with. I would also like to thank all the juniors for the amazing moments that we passed together.

Special thanks to Vahe and Marguerite. We three arrived at the same time, same day and shared a lot of great things. It was not always easy, but it is very important to have friends you can count on, whenever you need support or just a person to complain to. Vahe it will be weird not to see “Break” at my calendar. We started a nice tradition. Thank you for all the great moments and all the fun we had. Marguerite thank you for always being there for me. I will never forget the moment we met, when you were saying to people who were talking to me in English “mais elle parle Français”!!!

One of my biggest thank you to bureau 50. I had never imagined I would be that lucky to be in such a great office. You people made me happy to come to work every day. Benjamin what would this office do or be without you? Thank you for everything that you have done for me. Even when you did not know me very well you were of enormous help to me and always had an answer to all my questions! Caroline (mama) I still remember the first day we met and I know that you also remember my first “French related question”! It was a great pleasure to have you next to me and to talk to you (and I think we were talking a looooot)! Caro (bébé) maybe you did 100.000 hours of break with Pascal, but you did a lot more with me. I will miss your smile and your magnifique mood, all day, every day! Raissa thank you for all the great discussions and all the “administrative” help. I really enjoyed our discussions and I enriched, thanks to you, my cultural horizons! Tutu, you stayed only for a while but you really left your mark at this office. Thank you for all the fun and the laughter! And of course “Seni severim canim” and “Seni yargilamiyorum”! Encore merci pour tous les cours de Français ! Mais s’il vous plait, n’oubliez pas tout ce que vous avez appris maintenant que je serai plus là pour vous dire “ J’ai une question de Français” !!!

Of course this project would not be that unique if it was not for my IDEAS colleagues. I am very thankful to have had the opportunity and the honor to have met all these amazing people! We had such a great time together, travelled in so many places and shared new experiences. Jose, Fabiola, Elvira, Julia and Eleni thank you for these three years, but most importantly thank you for being there for me when I needed you. I never thought that I would leave this project with such valuable friends. Nico, Pavel, Enya, Haiyan,

Marius, Johanna, Saswati, and Arsenio thank you for everything!

Δεν θα γινόταν να κλείσω τις ευχαριστίες χωρίς να απευθύνω το μεγαλύτερο μου ευχαριστώ στην οικογένειά μου. Στην μητέρα μου, στον πατέρα μου και στον αδερφό μου. Ήταν είναι και θα είναι το μεγαλύτερο μου στήριγμα σε ότι και αν κάνω. Ήταν πάντα εκεί αυτά τα τρία χρόνια και με στήριξαν με όποιο τρόπο μπορούσαν και ξέρω ότι αυτό το διδακτορικό τους έκανε πολύ χαρούμενους και περήφανους. Σας ευχαριστώ πολύ για όλα!

And maybe for the last time...

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 633567.



Résumé

L'objectif de cette thèse était de développer des designs adaptatifs applicables dans les essais de phase I / II des thérapies moléculaires ciblées (TMC) en cancérologie. Les essais de phase I sont la pierre angulaire du développement de médicaments et leur correct design peut aider à améliorer la procédure globale. Une étape majeure pour améliorer les essais de phase I consiste à intégrer toutes les informations recueillies au cours de l'essai. Tout d'abord, nous nous sommes intéressés au développement d'un design, pour identifier la dose optimale (DO) qui peut prendre en compte plusieurs cycles de traitement et pas uniquement le premier, typique de ces essais. Le but principal était de prendre en compte les toxicités tardives et cumulatives. En outre, l'utilisation constante de biomarqueurs qui mesurent l'activité du médicament nécessite des approches alternatives, qui considèrent le biomarqueur comme une variable continue, mesurée à plusieurs moments au cours de l'essai. Il est bien connu que la dichotomisation de ces critères principaux ne peut que conduire à une perte substantielle d'information. De plus, considérer qu'une seule mesure de biomarqueur n'est pas suffisant pour définir l'activité du médicament, puisque l'évolution du biomarqueur change au cours du temps, indiquant une première réponse au traitement et éventuellement la progression de la maladie. Une autre question importante qui intéressait cette thèse était celle des données censurées. Lorsque la sélection de la DO est définie sur une longue période, c'est-à-dire 3-6 cycles, rencontrer des données manquantes est inévitable.

Pour la sélection de la DO, nous avons proposé un design adaptatif qui combine à la fois la toxicité et les données d'activité, mesurées sur une période de six cycles de traitement. Nous avons implémenté une technique de modélisation conjointe récente, qui prend en compte le temps avant la première toxicité dose-limitante (DLT) et de mesures répétées de biomarqueurs continus, sous un effet aléatoire partagé. Cette méthode de modélisation conjointe était basée sur l'inférence de vraisemblance et les paramètres du modèle ont été obtenus à partir de la maximisation numérique de la vraisemblance. La dose maximale tolérée (DMT) a été associée à un certain risque cumulé de DLT sur un nombre prédéfini de cycles de traitement. La DO a été définie comme la dose la moins

toxique parmi les doses actives, sous la contrainte de ne pas dépasser la DMT.

Pour étudier la performance du design et sa capacité à identifier correctement la DO, nous avons évalué un large éventail de scénarios. Plus précisément, nous avons envisagé de réduire la taille de l'échantillon, d'augmenter les variances, d'altérer la distribution des effets aléatoires, de supposer un risque croissant par cycle de traitement et finalement nous avons généré plusieurs scénarios sous divers degrés d'erreur de spécification des modèles. La méthode était très efficace en ce qui concerne la correcte sélection de la DO. Même en cas de censure excessive, avec seulement 7% des patients arrivant à la fin de l'essai, la méthode était robuste et les estimations très peu biaisées. En outre, notre design était assez fiable, puisque les participants de l'étude n'étaient généralement pas exposés à des doses hautement toxiques ou sous-thérapeutiques. Il s'agit du première design pour les essais de phase I qui comprenait plusieurs cycles de traitement, continus et répétés mesures de biomarqueurs, une définition cumulative pour la DMT et qui a abordé la question des réponses manquantes.

L'étude de la performance du nouveau design s'est focalisée sur le même exemple motivant de l'essai du cancer de l'ovaire afin de permettre la comparaison des deux approches. Des simulations approfondies ont montré que le nouveau design était très efficace sous divers degrés d'erreur de spécification du modèle. Une propriété importante du modèle de plateau était que lorsque les données étaient générées à partir du modèle de log, le pourcentage de recommandation correct de la DO était presque identique à celui observé lors de l'analyse des données avec le modèle de log. Ceci était une indication de l'adaptabilité du modèle à différentes trajectoires de biomarqueurs. Enfin, nous avons abordé le cas de l'échec de la convergence. Dans le modèle de log, certaines simulations n'ont pas pu être réalisées avec la modélisation conjointe, alors qu'avec le modèle de plateau et pour les mêmes simulations, aucun problème de ce type n'a été observé.

Une limite des deux designs adaptatifs à été liée à leur manque de comparabilité avec les méthodes existantes. Plusieurs méthodes de recherche de dose ont été proposées dans la littérature. Il existe des méthodes qui prennent en compte une période de DLT plus longue, sans toutefois prendre en compte plusieurs cycles de traitement ou une définition cumulative pour le DMT. Il existe également des méthodes qui considèrent l'activité

comme un critère d'évaluation continu, mais avec une seule mesure par patient. Par conséquent, la comparaison de notre design avec des approches alternatives n'était pas réalisable, car il ne serait pas possible de créer des scénarios correspondants.

La dernière partie de cette thèse a porté sur l'analyse de 27 études des TMCs de phase I, en tant que monothérapie. Ces essais ont été réalisés par le NCI et ont été fournis à l'EORTC et au DLT-TARGETT. Auparavant, nous avons discuté de l'importance de prendre en compte le temps lors d'un essai de phase I. Cependant, jusqu'à présent, il n'y avait rien dans la littérature, documentant le risque de toxicité sévère, pour plus d'un cycle de traitement. Par conséquent, notre intérêt principal était d'estimer le risque par cycle, ainsi que l'incidence cumulative de la toxicité sévère, jusqu'à six cycles de traitement. Un objectif secondaire était d'estimer les mêmes quantités, mais séparément pour les toxicités sévères hématologiques et non-hématologiques. Ces types de toxicités ne surviennent pas de la même manière et ils n'ont pas le même impact sur le corps humain, d'où l'intérêt d'estimer le risque séparément.

À cette fin, nous avons mis en œuvre un modèle de probit qui prend en compte le temps avant la première toxicité sévère et qui permet d'estimer à la fois le risque par cycle et l'incidence cumulative. Nous avons effectué l'analyse, d'abord séparément pour les groupes de patients traités aux doses ci-dessous, ci-dessus, et à la DMT, et ensuite sur tous les patients et nous avons fourni un nomogramme. Nous avons étudié le risque de toxicité pour les toxicités hématologiques et non-hématologiques, et enfin, nous avons effectué une analyse de sensibilité. Nous avons montré que l'incidence cumulative de la toxicité sévère chez les patients traités à la DMT était de 27% à la fin du premier cycle et elle a atteint 53% à la fin du sixième cycle. Pour les patients ayant reçu la DMT, l'incidence cumulative de la toxicité non-hématologique et hématologique était de 35% et 18%, respectivement.

Ce projet a été le premier à fournir des résultats sur le risque par cycle et l'incidence cumulative de la toxicité sévère sur plusieurs cycles de traitement. Nous avons également fourni des guidelines pour la bonne conduite des essais de phase I et des niveaux de toxicité cibles raisonnables, basés sur les résultats du DLT-TARGETT. Une limite de cette analyse est liée à la définition de toxicité sévère. Traditionnellement, la DMT est

définie sur la base des DLT survenant au cours du premier cycle de traitement. Les DLTs sont définies pour certaines toxicités de grade 3 et 4, et elles sont spécifiques à l'étude. Nous nous sommes intéressés au risque de toxicité sur plusieurs cycles, et pour cela nous aurions besoin de reconstruire la variable de la DLT pour les six cycles. Cependant, nous n'avons pas eu accès aux 27 protocoles d'étude. En conséquence, pour l'analyse, nous nous sommes concentrés sur la première toxicité sévère, qui ne coïncide pas nécessairement avec une DLT. Par conséquent, en considérant la première toxicité sévère, on aurait pu surestimer le risque de toxicité.

Avec cette thèse, nous avons voulu apporter la preuve qu'il y a de la place pour l'amélioration et à cette fin, nous avons développé un design adaptatif et une extension de ce design qui pourrait être appliqué à ces essais. Finalement, après l'analyse du DLTTARGETT, il est apparu que pour la définition de la dose recommandée pour les études de phase II, il est vital de considérer 3 à 6 cycles de traitement, d'envisager une définition cumulative du niveau de toxicité ciblé et, si possible, d'intégrer les mesures d'activité. En conclusion, l'intégration d'un plus grand nombre d'informations pourrait conduire à des essais de thérapies ciblées de phase I qui sont plus efficaces.

List of publications

- M. - A. Altzerinakou and X. Paoletti (2018). An adaptive design for the identification of the optimal dose using joint modeling of continuous repeated biomarker measurements and time-to-toxicity in phase I/II clinical trials in oncology. *Statistical Methods in Medical Research* (under revision).
- M. - A. Altzerinakou, L. Collette and X. Paoletti (2018). Cumulative Risk of Toxicity in Phase I Trials of Targeted Therapies: What to Expect at the Recommended Phase II Dose? *Journal of National Cancer Institute* (under revision).
- M. - A. Altzerinakou and X. Paoletti (2018). Optimal dose selection considering both toxicity and activity data; plateau detection for molecularly targeted agents (will be submitted).

Oral presentations

- Presentation of the thesis progression “Statistical methods for phase I/II trials of molecularly targeted agents in oncology”, IDEAS Network Meeting 21/06/2016, Traunkirchen, Austria.
- Presentation of the thesis progression “Statistical methods for phase I/II trials of molecularly targeted agents in oncology”, Lancaster Summer School-IDEAS Network Meeting 12/09/2016, Lancaster, United Kingdom.
- Presentation of the thesis progression “Statistical methods for phase I/II trials of molecularly targeted agents in oncology”, Institut Gustave Roussy 19/10/2016, Villejuif, France.
- Presentation of the thesis progression “Statistical methods for phase I/II trials of molecularly targeted agents in oncology”, Mid-Term Review European Committee IDEAS Network Meeting 31/01/2017, Brussels, Belgium.

- “An adaptive design for the identification of the optimal dose using joint modelling of efficacy and toxicity in phase I/II clinical trials of molecularly targeted agents”, Society for Clinical Trials 09/05/2017, Liverpool, United Kingdom.
- “An adaptive design for the identification of the optimal dose using joint modelling of biomarker measurements and toxicity over all treatment cycles in phase I/II trials of molecularly targeted agents in oncology”, Bayer Pharmaceutical company 30/03/2018, Berlin, Germany.
- “An adaptive design for the identification of the optimal dose using joint modelling of biomarker measurements and toxicity over all treatment cycles in phase I/II trials of molecularly targeted agents in oncology”, Applied Statistics Symposium 28/06/2017, Chicago, Illinois, United States of America.
- Presentation of the thesis progression “Statistical methods for phase I/II trials of molecularly targeted agents in oncology”, Torino Summer School-IDEAS Network Meeting 12/09/2017, Torino, Italy.
- “An adaptive design for the identification of the optimal dose using joint modelling of biomarker measurements and toxicity over all treatment cycles in phase I/II trials of molecularly targeted agents in oncology”, Janssen Pharmaceutical company 15/03/2018, Beerse, Belgium.
- “Phase I/II biomarker trajectory taking into account attrition due to excessive toxicity in the context of joint modeling”, CESP “Analyses de trajectoires” 27/03/2018, Villejuif, France.
- “Impact of cumulative risk of toxicity in phase I trials of molecularly targeted agents: What should we expect at the recommended phase II dose?”, Institut Gustave Roussy 11/06/2018, Villejuif, France.
- “Optimal dose selection considering both toxicity and activity data; plateau detection for molecularly targeted agents”, Joint Statistical Meeting 28/07/2018, Vancouver, Canada.

- “Optimal dose selection considering both toxicity and activity data; plateau detection for molecularly targeted agents”, Ideas Dissemination Workshop 26/09/2018, Novartis, Basel, Switzerland.

Poster presentation

- “Cumulative risk of toxicity in phase I trials of targeted therapies: What to expect at the recommended phase II dose?”, EFSPi Meeting 24/09/2018, Basel, Switzerland.

List of acronyms

ACE	Angiotensin Converting Enzyme
AE	Adverse Event
AIC	Akaike Information Criterion
AUSC	Areas Under Survival Curves
CA 125	Cancer Antigen 125
CIF	Cumulative Incidence Function
CRM	Continual Reassessment Method
CTEP	Cancer Therapy Evaluation Program
DLT	Dose Limiting Toxicity
EMA	European Medicine Agency
EORTC	European Organisation for Research and Treatment of Cancer
EWOC	Escalation with Overdose Control
FDA	Food and Drug Administration
ICS	Interval-Censored Survival
MAD	Maximum Administered Dose
MAR	Missing at Random
MaxED	Maximum Effective Dose
MinED	Minimum Effective Dose
MTAs	Molecularly Targeted Agents
MTD	Maximum Tolerated Dose
mTPI	Modified Toxicity Probability Interval

List of acronyms (cont')

NCI	National Cancer Institute
OBD	Optimal Biological Dose
OD	Optimal Dose
PSA	Prostate Specific Antigen
RP2D	Recommended Phase II Dose
TITE-CRM	Time-to-Event Continual Reassessment Method
TTL	Target Toxicity Level

Contents

Contents	xix
1 Dose finding in anti-cancer drugs	1
1.1 Thesis objective	1
1.2 Drug development	7
1.2.1 Phase II clinical trials	8
1.2.2 Phase III clinical trials	9
1.2.3 Phase IV clinical trials	10
1.3 Dose finding designs: Rule based designs	11
1.3.1 ‘3+3’ design	11
1.3.2 Other rule based designs	13
1.4 Model based designs	14
1.4.1 Continual reassessment method	14
1.4.1.1 Bayesian approach	15
1.4.1.2 Likelihood approach	19
1.4.1.3 Time-to-event CRM	21
1.4.1.4 Interval-censored survival model	23
1.4.1.5 Summary of CRM designs	24
1.4.2 Other model based designs	25
1.5 Molecularly targeted agents	26
1.5.1 Mechanisms of action	27
1.5.2 Designs combining toxicity and activity	29
1.5.2.1 Binary toxicity and activity	30

1.5.2.2	Time-to-event outcomes	37
1.5.2.3	Binary toxicity and continuous activity	39
1.5.3	Future concerns	42
2	An adaptive design for the identification of the optimal dose using joint modeling of continuous repeated biomarker measurements and time-to-toxicity	45
2.1	Introduction	45
2.2	Objective	46
2.2.1	Motivating example	46
2.3	Methods	48
2.3.1	General framework	48
2.3.2	Modeling time-to-DLT	49
2.3.3	Modeling activity	50
2.3.4	Joint modeling	51
2.3.5	Study design	56
2.3.6	Decision process	58
2.3.7	Estimands	61
2.4	Simulation settings and evaluation criteria	61
2.4.1	Framework	61
2.4.2	Time-to-DLT	62
2.4.3	Activity	62
2.4.4	Missing data	63
2.4.5	Main analysis	63
2.4.6	Sensitivity analysis	64
2.5	Results	66
2.5.1	Main results	66
2.5.2	Sensitivity analysis results	68
2.6	Discussion	69

3	Optimal dose selection considering both toxicity and activity data; plateau detection for molecularly targeted agents	73
3.1	Introduction	73
3.2	Methods	76
3.2.1	General framework	76
3.2.2	Modeling time-to-DLT	76
3.2.3	Modeling activity	77
3.2.4	Study design	79
3.2.5	Decision process	80
3.2.6	Independent modeling	83
3.3	Simulation settings and evaluation criteria	83
3.3.1	Framework	83
3.3.2	Main analysis	85
3.3.3	Sensitivity analysis	85
3.3.4	Comparison of joint modeling and independent modeling	86
3.4	Results	87
3.4.1	Main results	87
3.4.2	Sensitivity analysis results	87
3.4.3	Comparison of joint modeling and independent modeling	91
3.5	Discussion	92
4	Impact of cumulative toxicity in phase I trials. What should we expect at the recommended phase II dose?	97
4.1	Introduction	97
4.2	Objective	98
4.3	Methods	99
4.3.1	Trials and patients characteristics	99
4.3.2	Toxicity data	100
4.3.3	Main statistical analysis	100
4.3.4	Sensitivity analysis	103
4.4	Results	104

4.4.1	Trial characteristics	104
4.4.2	Treatment administration	104
4.4.3	Toxicity outcomes	104
4.4.4	Time-to-first severe toxicity	107
4.4.5	CIF of hematologic and non-hematologic severe toxicity	110
4.4.6	Sensitivity analysis	112
4.5	Discussion	114
5	Conclusion	119
	Bibliography	127
	Appendices	136
A		137
B		145
C		149
C.1	Descriptives	149
C.2	Main analysis	151
C.3	Complete model	153
C.4	Residual Check	154
C.4.1	For cycle effect	154
C.4.2	For dose effect	156
C.5	Competing events	158
C.5.1	Naive analysis	158
C.5.2	Competing risks analysis	160
C.6	Sensitivity analysis	162
C.6.1	Reclassification of mixed cases with non-hematologic toxicities	162
C.6.2	Exclusion of grade 3 hematologic toxicity	164

Chapter 1

Dose finding in anti-cancer drugs

1.1 Thesis objective

Before a new drug is approved as a medical product, it has to pass through a series of tests and trials in order to establish that it is both safe and active. The drug development process starts in the laboratories, *in vitro*, where they mainly test the new agent on cells, tissues or organs. After the agent has demonstrated some pharmacological activity, then further development proceeds with testing the agent on animals, also known as pre-clinical studies. The purpose of these tests is to examine the safety and activity of the new drug on animals and to suggest a starting dose or a range of doses, that will be later administered to humans. After the agent has been successfully tested on animals, we proceed with the *in human* clinical trials. These trials are further divided into four phases, called phase I, phase II, phase III and phase IV clinical trials. In this thesis we will focus on phase I and phase II trials.

The main objective in anti-cancer early phase trials is to establish a dose of a pharmacological agent (drug candidate) is safe and active. When a treatment regimen is considered to be safe that translates into not exceeding a predefined probability of toxicity. Activity is achieved when the agent has a “positive” treatment effect on the disease, i.e. partial or complete response.

More precisely, the principal goal of a phase I trial is to evaluate and establish the maximum tolerated dose (MTD) and the recommended phase II dose (RP2D) of the new

drug. Interest lies in the drug's safety and therefore, the primary endpoint is toxicity. Toxicity is categorized into 5 grades, with 1 showing minimum toxicity, whereas 5 indicates toxic death. Particularly, we are interested in the dose limiting toxicity (DLT), which is defined as severe toxicity (typically grade 3 or 4 toxicity) that entails dose reduction or treatment interruption (Table 1.1). The probability of DLT occurrence is assumed to monotonically increase with increasing dose, as shown in Figure 1.1 a). The MTD is the dose associated with some predefined probability of DLT, known as target toxicity level (TTL). The TTL is determined by clinicians before trial initiation and is typically between 20% and 33%. The RP2D is the dose administered to phase II trials and quite often it coincides with the MTD.

Phase I trials consist of small sample sizes, usually between 20 and 100 patients. In most diseases, phase I trials are conducted with healthy volunteers. However, in the circumstance of anti-cancer drugs volunteers are patients in the final stages of their disease. The study begins with the introduction of the first cohort that is a group of one or maximum three patients. These patients are assigned to a specific dose level, usually the first one, and they are followed for a pre-specified period of time. This period of time is divided into treatment cycles, typically 3-4 weeks. At the end of the first treatment cycle patients are evaluated as to whether they had a DLT or not. DLTs are defined solely for the first treatment cycle. The information collected from those patients is utilized in order to assign the next cohort of patients to the same or another dose level. Phase I trials are sequential designs, because patients enter the study in cohorts and not all at the same time. The motivation behind this is to impose a safety constraint in case of excessive toxicity. After all patients have been followed, for one treatment cycle, the MTD is defined as the dose producing a probability of DLT closest to the TTL. Therefore, dose finding is conventionally based on the toxicity data collected during the first treatment cycle.

This design for phase I trials dates from the era of cytotoxic agents and have been widely used over the years. Nevertheless, the emergence of molecularly targeted agents (MTAs) has challenged this type of design, due to their different mechanisms of actions. MTAs are often administered in chronic schedules, until DLT occurrence, disease progres-

Table 1.1: Table of definitions

Term	Definition
Adaptive Design	A trial design that allows modifications to some aspects of the design, based on collected data, after its initiation and without undermining the validity and integrity of the trial.
Adverse Event (AE)	Any untoward medical occurrence in a patient or clinical investigation subject administered a pharmaceutical product and which does not necessarily have to have a causal relationship with this treatment. †
Consent Withdrawal	A subject's decision to cease participation in a trial.
Dose Limiting Toxicity (DLT)	It is defined as the occurrence of a severe toxicity (typically grade 3 or 4 toxicity) that entails dose reduction or interruption.
Drug Activity	A measure of the physiological response that a drug produces. A less active drug produces less response, and a more active drug produces more response.
Expansion Cohort	Prolongation of phase I studies, with the addition of patients, often with different eligibility criteria, so as to better characterize the toxicity and activity profiles of experimental agents.
Maximum Effective Dose (MaxED)	The dose above which there is no clinically significant increase in pharmacological effect or activity.
Maximum Tolerated Dose (MTD)	The maximal dose acceptably tolerated by a patient population. It is usually the dose recommended for phase II trials.
Minimum Effective Dose (MinED)	The lowest dose producing a clinically important response different from the placebo response.
Pharmacodynamics (PD)	The study of the biochemical and physiological effects of drugs on the body.
Recommended Phase II Dose (RP2D)	The dose administered to patients in a phase II clinical trial, as that was chosen from phase I trial (Maximum Tolerated Dose).
Target Toxicity Level (TTL)	The maximum probability of DLTs that is considered acceptable in the trial. The TTL is used to select the MTD and is typically between 20% and 33% at the end of the first treatment cycle. † †
Treatment Cycle	In medicine, a course of treatment that is repeated on a regular schedule with possibly periods of rest in between. For example, treatment given for one week followed by three weeks of rest is one treatment cycle.

† ICH (1994)

† † Tourneau et al. (2009)

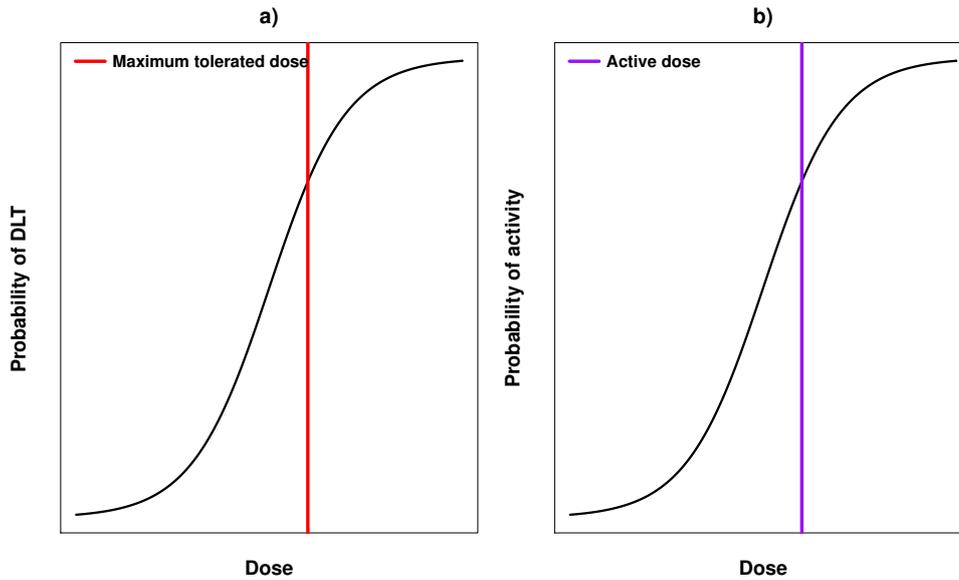


Figure 1.1: a) Dose-DLT relation. b) Dose-activity relation.

sion or consent withdrawal, rather than a predetermined number of cycles. Prolonged administration of these agents may induce late-onset or cumulative toxicities. Therefore, considering DLTs only from the first cycle of therapy can result in incorrect dose selection (Soria, 2011). In particular, Postel-Vinay et al. (2011) considered the data of 34 phase I trials of MTAs, from two large hospitals, with a total of 445 patients included. They showed a 20% incidence of grade 3 and 4 toxicities during the first treatment cycle with that number halving for each successive cycle. These results are pretty much in line with Postel-Vinay et al. (2014) and the analysis of the DLT-TARGETT database, which is composed of 54 completed phase I trials of MTAs as monotherapy. These discoveries clearly lead us to the conclusion that more attention should be paid to delayed toxicities, indicating that prolongation of the DLT period is crucial. Recently the European Medicines Agency (EMA) followed the report from the European Organization for Research and Treatment of Cancer -led DLT-TARGETT group (Postel-Vinay et al., 2014) and released a draft Guideline on the evaluation of anticancer medicinal products in man (European Medicines Agency, 2016). EMA stated that “in contrast to cytotoxic chemotherapy, MTAs are typically administered continuously and the toxicity profiles tend to differ so that DLTs may occur after multiple cycles of therapy. This is of im-

portance for the RP2D in cases where tolerability and toxicity guide dose selection, and may require alternative strategies with regard to definition of DLT and MTD”. The same guideline then recommends “Broader DLT definitions with longer DLT observation periods may therefore be relevant to consider. A distinction between cycle 1 acute toxicity, prolonged toxicity impacting on tolerability and late severe toxicity may be informative. Adverse events (AEs) should therefore always be reported by treatment cycle and the RP2D should be based on an integrated assessment of likely adverse reactions”.

Another challenge faced with MTAs is related to the dose-activity relationship. For cytotoxic agents the underlying premise is that both toxicity and activity increase monotonically with increasing dose (Figure 1.1 b)). The dose-toxicity relationship can serve as a surrogate for activity and thus there is merit in pushing the dose as high as safety allows. Consequently, the MTD is also the RP2D, as it is safe and at the same time the most active dose. Nevertheless, for MTAs this assumption may not hold (Kummar et al., 2006); it has been shown that the dose-activity curve of MTAs may reach a plateau after a certain dose level, countermining the widely accepted principle of the more the better. For that reason, EMA stated that “The use of pharmacodynamic endpoints (Table 1.1), where available, may also assist in dose selection”. For example, imagine a dose-activity relationship that plateaus after a certain dose level, as depicted in Figure 1.2. The red line represents the MTD. We can see from this figure that all doses located on the plateau have equal probability of activity. So the question raised here is what dose is the most appropriate to continue to subsequent phases. Is it possible to recommend a dose that is less toxic than the MTD, but equally active? Given the toxicities observed in a phase I trial in oncology, a lower dose with the same activity properties would constitute an optimal dose (OD).

From the above, it becomes clear that dose finding in the era of MTAs requires alternative methods that take into account both toxicity and activity, as well as treatment time (Reitsma et al., 2015). Phase I trials are a vital part of the drug development procedure, since inappropriate dose selection comes with a set of major drawbacks. To begin with, patients might be exposed to highly toxic or subtherapeutic doses, that cause undesirable adverse events or even death. Contrary to other diseases, in oncology

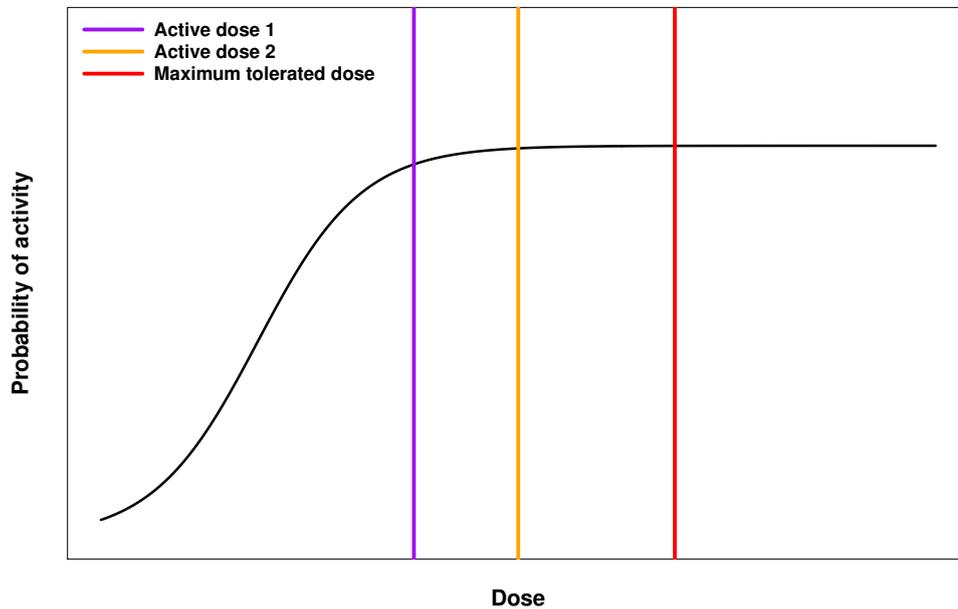


Figure 1.2: Dose-activity relation for MTAs.

exposure to low doses is considered an important issue. The reason is that we treat patients and not healthy volunteers and as a corollary under-treating them is equally dangerous as over-treating them. Another consequence of wrong dose selection after a phase I trial is the continuation or not to the next phases. In other words, a misspecified dose can result to either stop a potentially interesting drug or to promote a rather non-active or even hazardous one to the next levels. All mentioned above apart from an enormous cost on patients' life, they additionally have a colossal cost on the economy. Every year billions of dollars are spent for drug development, a big amount of which could be avoided if phase I and phase II trials improved. It is estimated that the cost of developing a new drug that will ultimately arrive in the market is 648 million US dollars (Prasad and Mailankody, 2017). What is more, millions of dollars are spent on agents that do not take approval to go into the market even though they succeeded to pass from phase I and II. Overall, the probability of success moving an anti-cancer drug from Phase I to the market is 3.4% (Wong et al., 2018).

The thesis is organized as follows. The remainder of this chapter is dedicated to

a review of several of the existing dose finding methods. We first present the two large categories of dose finding methods (rule based versus model based designs) and we discuss some of them in detail. Following, we talk about MTAs and their mechanisms of action. We introduce certain designs that consider both toxicity and activity for the OD selection. We close Chapter 1 with certain concerns regarding dose finding methods that need to be taken into account and addressed when designing a phase I trial. In Chapter 2, we present our adaptive design for the identification of the OD. We implement a joint modeling of continuous biomarker measurements and time-to-first DLT, using a shared random effect. We thoroughly explain the underline models, the study design as well as the decision process. Simulation scenarios with the corresponding results, follow, and we also discuss certain limitations of the proposed design, concerning model convergence and model flexibility under the case of a dose-activity plateau. Chapter 3 is an extension of our method presented in Chapter 2. In this extension the aim was to improve the dose-activity relationship, by allowing for more flexibility in the fitted model and improve arising issues regarding model convergence. In Chapter 4, we discuss the impact of cumulative toxicity in phase I trials of MTAs. More precisely, we analyzed toxicity data of 27 phase I trials of MTAs as monotherapy, with the objective to estimate the risk of toxicity and the cumulative incidence function per treatment cycle. Finally, we provided recommendations regarding the design and conduct of phase I trials. Conclusion and discussion of the thesis follow in Chapter 5.

1.2 Drug development

Before proceeding with the dose finding designs we will briefly introduce phase II, phase III and phase IV clinical trials of the drug development.

1.2.1 Phase II clinical trials

Once a phase I trial is completed, we continue with phase II clinical trials, also exploratory studies. Phase II are called “Proof of Concept Studies”, since their main goal is to assess the activity of the new drug. Consequently, activity constitutes the primary end point. Phase II enroll 50 to 300 patients and they last for a few years. Eligibility criteria are stricter than in phase I trials, with participants having a specific type of cancer. It is very often required that the patients’ condition is not too critical or too mild, so as to better investigate the activity. In these trials the RP2D is tested with the purpose to define its activity. It is possible that the drug is compared to a control compound, even though this is not always the case. Finally, we wish to obtain more information with regard to the pharmacokinetics and pharmacodynamics properties of the agent.

Outside of oncology, certain parameters of interest within phase II trials are the minimum and the maximum effective dose (MinED),(MaxED), as shown in Figure 1.3. The minimum effective dose is defined as the lowest dose producing a clinically important response that can be declared statistically significantly different from the control response, whereas the maximum effective dose is the one above which there is no clinically significant increase in pharmacological effect or efficacy. Both quantities can be of interest in clinical trials, since they can help to better understand and evaluate the properties of the drug under investigation. In oncology, since we investigate only one dose level, we do not consider MinED and MaxED, even though it would be of great interest.

For these studies to be considered successful both safety and activity must be shown. Then and only then it is safe to proceed with a phase III trial. Phase II trials, including successful ones, come with a set of certain limitations. First of all, due to the narrow patient population it is difficult to generalize the results to a broader one. What is more, due to their duration, their sample size and usually the lack of control group, phase II cannot provide definite and detailed conclusions.

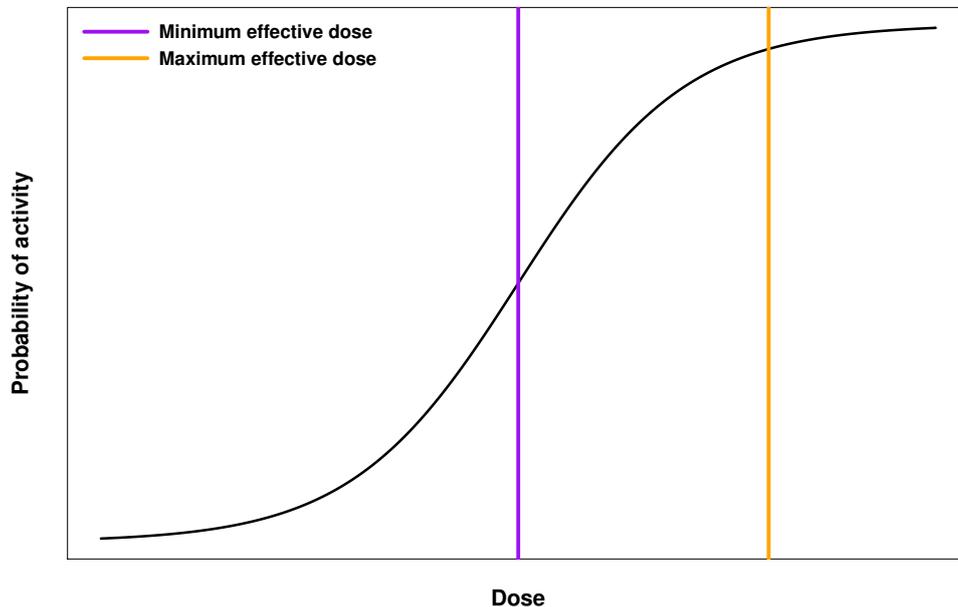


Figure 1.3: Dose-activity relation.

1.2.2 Phase III clinical trials

Phase III clinical trials are large confirmatory trials that aim to compare the new treatment to the standard of care and establish the results obtained from previous trials. The goal is to further investigate and prove the clinical activity of the regimen, compared to the standard of care and to evaluate the risk-benefit balance.

Phase III are large studies usually of several hundreds to thousands patients and can last up to several years. The eligibility criteria for patient inclusion is less restricted than phase II and as a result generalizability of the findings can be accomplished. In these trials there is always a control group of patients, to whom another drug is administered. This latter is known as standard treatment and it has already been shown to work. If a standard treatment does not exist, then the new agent is compared to placebo.

Once participants are enrolled into the study, they are randomized between the experimental and the standard treatment. Randomization ensures that possible confounders are eliminated and the two or more groups are comparable. Another important aspect of phase III is blinding. Patients or both physicians and patients, if possible, should not be

aware of the treatment administered.

Even though recommendation of doses is part of early phase trials, it is still possible that dose selection is further refined in phase III. One reason for this is that because of study duration patients are exposed to the drug for a longer period of time and consequently later on or cumulative toxicities can be recorded and taken into account. An example is the re-evaluation of the MTD of MTAs, in phase II and in phase III trials (Le Tourneau et al., 2010b; Iasonos et al., 2012; Jardim et al., 2014).

After a phase III is completed a study report is transmitted to the European Medicines Agency (EMA) or the Food and Drug Administration (FDA) for United States and is evaluated. These regulatory agencies are responsible to decide whether the new agent can go to the market for public use or get rejected (Crowley and Hoering, 2012).

1.2.3 Phase IV clinical trials

After the successful completion of phase I, II and III clinical trials and the approval by the regulatory agencies, the drug arrives to the market. Evaluation of the drug can continue even when it is on the market. These studies are known as phase IV or post-marketing studies.

One major advantage of these studies is that they account for the full patient population. As a result, it is feasible to record rare adverse events that were not detected during the previous trials. What is more, post-marketing surveillance comes with a large scale duration compared to the past trials. This way, cumulative effects of the concentration of the drug can also be detected and dealt with. Phase IV designs can use various designs, such as a randomized trials with control arm, but single-arm, non-randomized or open-label studies are also accepted. Phase IV trials in oncology aim as well to investigate individualized or personalized treatment and thus, randomization or blinding is not that often (Zhang et al., 2016).

Compared to premarketing phase I–III trials, phase IV studies evaluate drug safety in a real-world setting, which may provide evidence to ensure or further refine the safety

and activity of approved drugs.

1.3 Dose finding designs: Rule based designs

Dose finding methods are divided into two categories, the rule based and the model based. Rule based, also known as “up and down” designs are widely utilized, because of their simplicity. Their principal aspect is that they only make one assumption; toxicity increases monotonically with dose. Dose escalation or de-escalation proceeds based on certain rules, that vary among these methods. In the next section, a few of them will be shortly presented. On the other hand, model based designs, they do assume a model that describes the dose-toxicity relationship and they can be more complicated to implement. A particular model based design will be analytically discussed in Section 1.4.

1.3.1 ‘3+3’ design

The ‘3+3’ design is probably one of the most commonly used designs since till now it has been applied on 95% of published phase I trials clinical trials (Ji and Wang, 2013). The basic idea behind this method is that participants enter the study in groups of three, also known as cohorts. The MTD is defined as the dose with a percentage of toxicity equal or below 33%. Every cohort is allocated to a specific dose and then is followed for one treatment cycle. Based on their toxicity outcome, i.e. the number of observed DLTs, the next cohort of patients is assigned to the same, the next or the previous dose level, if possible. The principal rule is that if non of the three patients had a DLT then the next cohort will be assigned to the next dose level. If one of the three had a DLT, then the next cohort will get the same dose. Finally, if two or more DLTs have occurred, then this dose is defined as the maximum administered dose (MAD). The MAD is the dose that is considered to produce a probability of toxicity greater than the MTD, so it will not be given to any more patients. In this case three more patients will be assigned to the

Table 1.2: The 3+3 rule based method

No. of patients with DLTs at a given dose level	Escalation decision rule
0 out of 3	Enter 3 patients at the next dose level.
1 out of 3	Enter at least 3 more patients at the same dose level.
≥ 2 out of 3	Dose escalation will be stopped. This dose will be declared the MAD. Enter 3 more patients at the next lowest dose level, if only 3 patients were treated previously at that dose.

previous level, if only three patients had been treated to that level before. Otherwise, the previous dose level is declared MTD. The trial halts after observing 2 or more DLTs at a given dose level and 1 or none DLTs for the 6 patients treated at the previous dose level. So we systematically identify the MTD or we stop enrollment after the first 3 or 6 patients due to excessive toxicity at the first dose level. Table 1.2 is given so as to better clarify how this method works.

As mentioned above, this technique has been widely applied over the years due to its simplicity. However, there are certain limitations that need to be taken into account when implementing this design. First limitation is the poor accuracy of the estimation of the MTD, since its identification can be achieved after only six patients have been enrolled. Furthermore, due to the size of the cohorts and because dose escalation begins from the first dose level, many patients are ultimately treated at low or else sub-therapeutic doses. Additionally to that, only few of the participants will eventually be assigned to the MTD or to doses close to the MTD. Given the severity of the patients' condition, under-treating them is not ethical, since it can prove to be fatal. Another drawback of the '3+3' is the lack of flexibility, increasing this way the risk of recommending a phase II dose, that is below the true MTD (Tourneau et al., 2009).

1.3.2 Other rule based designs

Several rule based designs can be found in bibliography, such as variations of the ‘3+3’ (‘2+4’, ‘3+3+3’, pharmacologically guided dose escalation) (Tourneau et al., 2009). Another variation of the ‘3+3’ that was proposed to shorten the timeline of pediatric phase I trials, when there is no prior information about the dose range under investigation, is the rolling six design. With this design, two to six patients are allowed to be concurrently enrolled in the study. Every cohort starts with three patients entering the study and being assigned to a dose level. If toxicity information is available for all three of them and no DLT has occurred then the fourth patient is assigned to the next dose. If data is not available for one or more of the first three patients or if one DLT has been observed then the fourth patient is administered to the same dose. Finally, if two or more DLTs have occurred the fourth participant is assigned to the previous dose. The same procedure is followed for patients number five and six (Skolnik et al., 2008; Eisenhauer et al., 2015). According to Skolnik et al. (2008), the two methods, the ‘3+3’ and the rolling six, have been shown to have similar operating characteristics. Additionally, Doussau et al. (2012) showed in simulation scenarios that the probability of identifying the correct dose is the same between the two designs. Accelerated titration design was developed to reduce the number of patients treated at subtherapeutic dose levels and the trial’s duration, as well as to collect important information, such as cumulative toxicity (Simon et al., 1997). To that end, three designs were created, with cohorts of one patient. If a patient experienced a DLT or two patients experienced grade 2 toxic effects, during the first cycle for designs 1 and 2 and during any cycle for design 3, then the design switched to the ‘3+3’. Accelerated titration design allows for inpatient dose escalation. The authors compared their models with the ‘3+3’ method and concluded that their method reduced the number of patients who were undertreated and accelerated the completion of the trial.

Rule based methods have been adopted to a great extent, mostly because they are easy to implement, with regard to statistical software. What is more, they seem to be more easily accepted by institutional review boards. Despite all this, these methods are insufficient in establishing a dose that corresponds to a certain target toxicity level. Furthermore, due to the way they function, they do not consider all information gathered

during a phase I. Consequently, in the past years implementation of model based designs has increased, even though these techniques exist for many years now.

1.4 Model based designs

1.4.1 Continual reassessment method

Model based designs are an alternative proposal for the dose finding process. The principal difference with rule based dose escalation techniques is that they utilize statistical models to describe the dose response relationship. Additionally, toxicity outcomes of all patients, observed during the trial, are used so as to better estimate the model parameters and allocate patients to appropriate doses. Statistical inference can be either Bayesian or frequentist. One very popular approach that constitutes the cornerstone of model based techniques is the continual reassessment method (CRM). The CRM and its variations are thoroughly discussed in the next sections, since they are of great importance for this project.

The CRM was introduced by O'Quigley et al. (1990) in 1990. The CRM can be summarized in 7 steps.

1. Select a TTL, based on clinician's views and a dose-toxicity function.
2. Enroll one or more patients and treat them at the first dose level or a dose chosen based on prior knowledge.
3. Follow the patients for a time period (e.g. for one treatment cycle) and record whether they have a DLT.
4. After the follow up period and before the entry of the next patients update knowledge on the dose-toxicity function.
5. Using either Bayesian methods or the classic likelihood approach estimate the new probability of toxicity at each dose level.

6. Finally, choose the dose level that has the probability of toxicity that is closest to the TTL.
7. Repeat this process before the entry of every new patient and till you reach the maximum sample size or study stops due to excessive toxicity.

The main objective is to suggest a dose with a given TTL and to concentrate experimentation at that which all current available evidence indicates to be the best estimate of this level. This technique is based on constantly updating our knowledge with regard to the dose-toxicity relationship, as observations become available. Patients are always treated at the dose whose response probability, based on the observations of all patients, is closest to the desired level.

In this section we present the original design suggested by O'Quigley et al. (1990), as well as alterations recommended over the years with the goal to improve the CRM. The approach by O'Quigley et al. (1990) was based on Bayesian statistics, whereas a maximum likelihood approach was later on suggested.

Consider a range of M planned doses to be tested $D = \{d_{(1)}, \dots, d_{(M)}\}$. Let Z_i be a binary random variable $(0, 1)$, with 1 indicating DLT for the i^{th} individual, $i = (1, \dots, N)$. We assume an increasing dose-toxicity function $E(Z_i)$ that is depicted by $\psi(d_i, a)$, where $a \in \mathcal{A}$ is a parameter. We assume that $\psi(d_i, a)$ is monotonic in d and a . Now, let ϵ be the TTL. Then, there is a d^* for which $\psi(d^*, a_0) = \epsilon$, where a_0 is the true state of nature.

1.4.1.1 Bayesian approach

Let $\Omega_i = ((z_1, d_1), \dots, (z_{i-1}, d_{i-1}))$ and suppose $f(a, \Omega_i)$ to be a positive function or else the prior, that contains all the information for a . Assuming that $\mathcal{A} = (0, \infty)$, then

$$\int_0^\infty f(a, \Omega_i) da = 1. \quad (1.1)$$

Here, $f(a, \Omega_i)$ is the prior before the entry of the i^{th} observation. Then, we want to estimate the posterior distribution $f(a, \Omega_{(i+1)})$. Taking Bayes' theorem, we need the

prior 1.1 and the model likelihood for subject i that is

$$\phi(d_i, z_i, a) = \psi^{z_i}(d_i, a) \{1 - \psi(d_i, a)\}^{(1-z_i)}.$$

Then, the posterior distribution is

$$f(a, \Omega_{(i+1)}) = \frac{\phi(d_i, z_i, a) f(a, \Omega_i)}{\int_0^\infty \phi(d_i, z_i, v) f(v, \Omega_i) dv}.$$

The probability of toxicity at dose level l after the enrollment of i patients is given by

$$p_{(l)} = \int_{\mathcal{A}} \psi(d_{(l)}, a) f(a, \Omega_i) da.$$

Instead of working with the expected values of the probabilities over \mathcal{A} the authors proposed to work directly with the expected value of a over \mathcal{A} . Thus,

$$\hat{p}_{(l)} = \psi\{d_{(l)}, \mu\} \quad \mu = \int_{\mathcal{A}} a f(a, \Omega_i) da.$$

After having estimated all the above quantities, the next patient is assigned to the dose that produces a probability of toxicity closer to the TTL.

For a better understanding, the method will be illustrated by an example from O'Quigley et al. (1990).

1. As discussed above, the first step is to define the TTL ϵ , here 20%. Suppose now that we want to test 6 parameterized dose levels $D = (-1.47, -1.1, -0.69, -0.42, 0.0, 0.42)$. The authors have chosen the hyperbolic tangent function to depict the dose-toxicity relationship

$$\psi(d_i, a) = \left\{ \frac{\tanh d_i + 1}{2} \right\}^a, \quad (1.2)$$

and for the prior distribution they proposed the exponential function

$$f(a, \Omega_1) = \exp^{-a}.$$

2. The first patient is allocated to the prior estimation of the MTD. In our function \hat{a} is assumed to be equal to 1, since this is the mean value of the exponential distribution that we have chosen as prior. The third dose level $d_{(3)} = -0.69$ gives a probability of DLT $\hat{p}_{(3)} = 20.1\%$, which is the closest to $\epsilon = 20\%$. Thus the first patient will be treated at dose level 3.

3. After the first patient is enrolled, he will be followed for one treatment cycle, so as to observe whether DLTs occur. Assuming that no DLT occurs, $z_1 = 0$, the next step is to estimate the dose to which the second patient will be allocated to. To do so we need to take into account the first patient's toxicity outcome to estimate the posterior distribution of a . For the posterior distribution we take the likelihood for the first patient and the prior distribution. For the likelihood,

$$\begin{aligned} \phi(d_{(3)1}, z_1, a) &= \psi^{z_1}(d_{(3)1}, a) \times \{1 - \psi(z_{(3)1}, a)\}^{(1-z_1)} \\ &= \left(\frac{\tanh d_{(3)1} + 1}{2}\right)^{a*0} \times \left(1 - \left(\frac{\tanh d_{(3)1} + 1}{2}\right)^a\right) \\ &= \left(1 - \left(\frac{\tanh d_{(3)1} + 1}{2}\right)^a\right). \end{aligned}$$

So the posterior distribution of a after the first patient is observed takes the form

$$f(a, \Omega_2) = \frac{\left(1 - \left(\frac{\tanh d_{(3)1} + 1}{2}\right)^a\right) \times \exp^{-a}}{\int_0^\infty \left(1 - \left(\frac{\tanh d_{(3)1} + 1}{2}\right)^v\right) \times \exp^{-v} dv}.$$

We proceed by calculating the mean of the posterior distribution as shown below,

$$\begin{aligned}
 \mu_2 &= \int_0^{\infty} a \times f(a, \Omega_1) \, da \\
 &= \frac{\int_0^{\infty} \left(1 - \left(\frac{\tanh d_{(3)} + 1}{2}\right)^a\right) \times \exp^{-a}}{\int_0^{\infty} \left(1 - \left(\frac{\tanh d_{(3)} + 1}{2}\right)^v\right) \times \exp^{-v} \, dv} \\
 &= 1.38.
 \end{aligned}$$

4. Having summarized the information of the prior knowledge and the one provided after the first patient, we use that information to estimate the posterior probability of DLT at every dose level. To do so, we substitute in the function $\psi(d_i, a)$, a with μ_2 and we calculate the $\hat{p}_{(l)}$ for every $d_{(l)} \in D$. For instance, for the first dose in the dose range

$$\begin{aligned}
 \hat{p}_{(1)} &= \left(\frac{\tanh d_{(1)} + 1}{2}\right)^{\mu_2} \\
 &= 0.0161104 \quad \text{or} \quad 1.6\%.
 \end{aligned}$$

5. After obtaining all $\hat{p}_{(l)}$ we select the one that gives a probability of DLT closer to the TTL, by absolute value i.e. $|\epsilon - \hat{p}_{(l)}|$. That dose is administered to the second patient. In this case that would be the fourth dose level $d_{(4)} = -0.42$ with a probability of DLT equal to 19.1%.
6. We repeat steps 3-5 before the entry of every new patient.

General remarks

The authors have run extensive simulations with the goal to test the operating characteristics of their design and to investigate how it performs under various scenarios.

Initially, O’Quigley et al. (1990) tested three different models. They concluded that all models seemed to perform quite well, in most of the scenarios tested, which was a proof of the robustness of the method. That means that the models were sufficiently successful in choosing the right dose level, and not treating many patients in subtherapeutic or highly toxic doses. However, later on Shen and O’Quigley (1996) pointed out the importance of working with a one-parameter model. They argued that a single parameter suffices to model the toxicity at any given level. Once the method has settled at some level the introduction of a second parameter poses the problem of instability of the parameter estimates. This issue does not arise when considering continuous doses because, even though we may have a single target, each subsequent experiment takes place at a distinct design level, thereby providing adequate information to fit both slope and intercept in a linear setting.

O’Quigley and Shen (1996) discussed a concern that had been the source of worrying when the CRM was originally introduced. Their concern was related the choice of prior (Mick and Ratain, 1993). In the case phase I trials, little or no information exists on the molecule under investigation. Even when the prior is vague enough, it can be argued that its selection can have an impact on the estimation of the posterior, under very small sample sizes. Despite the fact that dependence on the prior diminishes as the sample size increases, it is true that in phase I trials, and especially in the past, most of the times the sample size does not exceed 20 or 30 patients. Therefore, attention should be paid on the prior selection.

1.4.1.2 Likelihood approach

In an attempt to address the above concerns O’Quigley and Shen (1996) presented an alternative version of the CRM that is based on the maximum likelihood theory. This design is identical to the Bayesian approach, with the only difference being in the initiation of the escalation scheme. In particular, under the maximum likelihood, parameter estimation is not feasible before some heterogeneity is encountered in the responses. To maximize the likelihood, it should be non-monotone, so that the estimates are not on the boundary of the parameter space, i.e. $(0, \infty)$. That translates into having at least one

DLT and one non-DLT in the data, given the one-parameter model. This is accomplished by introducing some initial escalation scheme. This scheme can be a ‘3+3’ design or any other kind of “Up and Down” schemes. After some patients are treated and heterogeneity is observed, we can then switch to the CRM.

Taking the previous example, once again we select a TTL that is $\epsilon = 20\%$ and a range of 6 doses $D = (-1.47, -1.1, -0.69, -0.42, 0.0, 0.42)$. Contrary to the Bayesian design, the first cohort of patients is always treated at the first dose level, $d_{(1)} = -1.47$. The dose-toxicity relationship is depicted by a reparameterization of the hyperbolic tangent function 1.2 that is

$$\psi(d_i, a) = a_i^a, \quad \text{where } a_i = \left\{ \frac{\tanh d_i + 1}{2} \right\}.$$

Then, the log-likelihood takes the form

$$V(a) = \sum_{m=1}^{i-1} z_m \log \psi(d_m, a) + \sum_{m=1}^{i-1} (1 - z_m) \log(1 - \psi(d_m, a)).$$

After estimating the model parameter, we use that information to reassess the probability of toxicity at every dose level. We substitute in our function $\psi(d, a)$ the \hat{a} and we estimate the $\hat{p}_{(l)}$ for every $d_{(l)} \in D$. Similarly to the Bayesian approach, after obtaining all $\hat{p}_{(l)}$, we identify which dose level gives a probability of toxicity closer to the TTL, and that dose is recommended for the next patient.

General remarks

O’Quigley and Shen (1996) have run simulations with the purpose to compare the Bayesian and the maximum likelihood approach as far as dose allocation and final recommendation is concerned. Based on the simulations, the authors concluded that no significant differences were observed among the models regarding the recommended dose level at the end of the trial.

An advantage of this approach is that no prior is requested and calculations are

somehow simplified. What is more, researchers might feel more comfortable with this approach because no initial “guess” is required concerning the dose allocation of the first patient. However, the authors argued that the purpose of this method was not to improve any of the existing methods, but rather to suggest a new formulation that would seem more “friendly”.

1.4.1.3 Time-to-event CRM

One of the major challenges, with regard to phase I trials and the CRM, is the sequential nature of the design. Before the entry of a new cohort, it is required that all patients, already in the study, are fully observed. The main reason is that decisions on dose allocation are based on the toxicity outcomes observed from the previous patients. Consequently, a study could last for years before recruiting all patients. Additionally, the DLT is defined with respect to some time duration, that is usually one treatment cycle. Even though this does not pose a problem for diseases that progress rapidly, in oncology toxicities can occur after the first treatment cycle. Therefore, a DLT assessment period can last beyond one treatment cycle, since late-onset toxicities is a common phenomenon and they should be taken into account when estimating the MTD.

Cheung and Chappell (2000) proposed a method that embodies the time-to-event outcome of each patient into the CRM (TITE-CRM). The purpose of this technique is to decrease the total duration of a phase I trial and at the same time exploit the benefits of the sequential design. In other words, a patient can enter the study before the previous patients have completed a treatment cycle, a property that can significantly decrease the duration of the study. The basic notion of TITE-CRM is to estimate the dose toxicity curve based on the current toxicity status of all patients, including those who have not been followed up to the end of the DLT assessment period. That is accomplished by assigning weights to every patient in the likelihood.

Assuming again a dose-toxicity function that increases monotonically with dose and T to be the total duration of follow-up, called “observation window”. Each patient in the study carries a weight w that is $0 \leq w \leq 1$. Patients entering the study are assigned to a weight equal to 0 and the longer their follow-up the higher becomes the weight. A patient

can have a weight equal to 1 in two cases, either if they had a DLT or if they completed the follow-up. Finally, we denote by S_i the time to toxicity of the i^{th} individual. Then for $s \leq T$,

$$\begin{aligned} P(S_i \leq s) &= P(S_i \leq s | S_i \leq T)P(S_i \leq T) \\ &\equiv w(s; T)\psi(d_i, a). \end{aligned}$$

The likelihood function is based on conditionally independent current toxicity outcome data. The authors selected a linear weighting scheme that is described by the following function

$$w(s; T) = \min\left(\frac{s}{T}, 1\right).$$

Independently of the weighting formulae, the core of the CRM method remains the same. After selecting the weighting scheme, we estimate the model parameters and finally the probability of DLT at every dose level in order to select the dose for every new entry.

General remarks

Extensive simulations were conducted comparing three different types of CRM with the TITE-CRM. The maximum follow-up duration for every patient was 6 months and two different sample sizes were tested $n_1 = 25$ and $n_2 = 48$. The range of dose levels was equal to 6. The purpose of this process was to compare the accuracy of each method in selecting the correct final dose and the study duration.

CRM simulations gave results comparable to its TITE counterparts. As far as the final recommendation was concerned models seemed to produce quite similar results. It was the in-trial allocation that slightly changed depending on the time failure model that was chosen. In addition to this, the study duration in all cases of TITE was significantly shorter than the CRM counterparts. More simulation results can be found in the Cheung and Chappell (2000) article. It is worth mentioning that when delaying accrual is con-

sidered unethical, i.e. when patients should enter the study before the previous patients have completed their follow-up, the TITE-CRM can be a good solution in assigning patients to dose levels, given all available information. Even though this technique does not improve the accuracy of CRM, it affects in-trial allocation. What is more, according to Cheung (2011) it can be useful when late onset toxicity is prominent or we expect non negligible dropout. However, in the case of excessive dropout bias is induced in the estimations (Cheung, 2011).

1.4.1.4 Interval-censored survival model

Most dose finding designs estimate the probability of DLT after one treatment cycle. Consequently, all DLTs occurring after the first cycle are not considered discarding a big fraction of the information acquired during a phase I trial. Sinclair and Whitehead (2014) addressed this concern by proposing an interval-censored survival (ICS) model, which models the probability that the first occurrence of a DLT is in a particular treatment cycle, conditionally on having no DLTs in the previous cycles.

Let $C = \{1, \dots, k\}$ denote the treatment cycles and k be the last cycle in the trial. Then for the probability of DLT $p_{(l)s}$ in cycle s and at dose level l we have

$$p_{(l)s} = \begin{cases} (\pi_{(l)1}) & s = 1 \\ (1 - \pi_{(l)1})(1 - \pi_{(l)2}) \dots (1 - \pi_{(l)(s-1)})\pi_{(l)s} & s = 2, \dots, k \\ (1 - \pi_{(l)1})(1 - \pi_{(l)2}) \dots (1 - \pi_{(l)(s-1)})(1 - \pi_{(l)s}) & s > k \end{cases}$$

where $\pi_{(l)s}$ is the conditional probability of DLT in cycle s and at dose l . Let S be the probability of not having a DLT. For the dose-toxicity relationship Sinclair and Whitehead (2014) chose a linear model for the complementary log-log transformation of $\pi_{(l)s}$ so that

$$\begin{aligned} \log(-\log(1 - \pi_{(l)s})) &= \log\left(-\log\left(\frac{S > s}{S > s - 1}\right)\right) + a_2 \log d \\ &= a_s + a_2 \log d \end{aligned}$$

where a_s is the model intercept that is cycle specific and a_1 is the parameter for the dose effect. The dose assigned to the next patient is selected on the basis of the dose that has the modal posterior probability of a DLT closest to the TTL after k cycles of treatment. This distance is calculated via a gain function

$$\hat{g}^{(l)} = \frac{1}{(\epsilon - \hat{p}^{(l)k})^2}.$$

An advantage of the ICS model compared to the previously presented TITE-CRM is related to missing values. The ICS model allows for censoring and dropout, while at the same time taking advantage of all available information. According to Sinclair and Whitehead (2014) a limitation of their method is that it assumes a model for the data that may not be appropriate in some situations, for instance, in the case of a non-linear dose effect.

1.4.1.5 Summary of CRM designs

The CRM is a model based technique for the evaluation of the MTD in phase I clinical trials. It is widely accepted and utilized over the past few years. As a method it is accurate enough concerning the recommendation of the final dose and most importantly it exploits all the knowledge gathered throughout the trial. After extensive simulations O'Quigley et al. (1990) and O'Quigley and Chevret (1991) showed that the CRM outperformed other sequential designs (Neuenschwander et al., 2008; Berry et al., 2010; Cheung, 2011). More specifically, in the context of phase I cancer trials, rule based designs were inferior with regard to identifying the MTD and for that reason O'Quigley and Chevret (1991) recommended that the FDA and other regulatory agencies encourage the use of CRM in these trials. Shen and O'Quigley (1996) and later Cheung and Chappell (2002) showed that even under model misspecification asymptotic convergence of the CRM algorithm to the correct dose is feasible, under certain conditions. However, investigators need to be careful when selecting the dose-toxicity model and the number of included parameters. Paoletti and Kramar (2009) argued that under limited sample sizes, typical of phase I, asymptotic convergence can be affected by model choices. In their article, they examined

a set of scenarios for 4 working models, under both the Bayesian and the likelihood framework. Their principal conclusion was that model selection has a significant effect on the operating characteristics.

In this section, two different types of CRM algorithms were described; the Bayesian and the maximum likelihood. It is pointed out in several occasions that despite the fact that both approaches produce quite similar results, the likelihood approach tends to be more easily accepted, since it is more conservative and does not require prior information. As far as the TITE-CRM is concerned, shortening the study duration is an essential quality of this approach, even though investigators are required to be extra careful with its implementation.

1.4.2 Other model based designs

Escalation with overdose control (EWOC) is an alteration of the Bayesian CRM that was proposed by Babb et al. (1998) and Rogatko et al. (2005), with the goal to improve the safety of the CRM. “The only difference between the EWOC and the CRM is that with the EWOC method, the probability of administering a dose that exceeds the MTD for each higher dose level is assessed after each patient, with an interdiction of dose escalation if this probability exceeds some critical prespecified value” (Tourneau et al., 2009). Despite the fact that this technique can protect participants from being treated to highly toxic levels, it seems to be more conservative and it requires flexible models with the addition of more informative priors.

The last model based design presented here is the “modified toxicity probability interval” (mTPI) design, proposed by Ji and Wang in 2013. It is based on the principal idea of the ‘3+3’ design, however it is categorized as model based since it makes use of probability models. In the mTPI approach dose escalation or de-escalation proceeds with respect to toxicity probability intervals, created based on the TTL. Let us assume that the TTL is 20%. Then three probability intervals are constructed. The first one (underdosing) is [0%, 15%), the second one (proper dosing) is [15% , 26%) and the third

one (overdosing) is [26% , 100%]. In order to decide to what dose level each patient will be assigned to, the toxicity rate of the currently used dose is taken into account. If this rate is within the underdosing interval, then dose escalation is the way to proceed. If the rate is in the proper dosing range then continuing at the same dose level is recommended. Finally, de-escalation is suggested in case of overdosing. The mTPI is simple to implement and it is quite safe, since it has a lower probability of exposing patients to highly toxic doses, compared to the 3+3.

1.5 Molecularly targeted agents

Routinely administered anti-cancer treatments, such as chemotherapy, are called cytotoxic drugs. Cytotoxic agents are medicines used to treat cancer, because of their ability to kill cancer cells. They tend to work by interfering with some aspect of how the cells divide and multiply. For instance, some work by affecting the cells' genetic "makeup". There are many different cytotoxic medicines used in the treatment of cancer. In each case the one (or ones) chosen depend on the type and stage of cancer.

Nevertheless, in the past years the emergence of novel agents as a new therapeutic option has caused a major shift in the treatment of cancer. These agents are called molecularly targeted agents (MTAs) or targeted therapies. Targeted cancer therapies are drugs or other substances that target a protein or an enzyme that carries a mutation or other genetic alteration that is specific to cancer cells and not found in normal host tissue. Particularly, these proteins are involved in processes such as cell signaling, oncogenesis or tumor suppression, cell cycle regulation, angiogenesis, immunologic pathways, metastasis and apoptosis. Therapies that have been approved for treatment of cancer include hormone therapies, signal transduction inhibitors, gene expression modulators, apoptosis inducers, angiogenesis inhibitors, immunotherapies, and toxin delivery molecules. Cancer vaccines and gene therapy are sometimes considered targeted therapies because they interfere with the growth of specific cancer cells (National Cancer Institute, 2018).

1.5.1 Mechanisms of action

In order for MTAs to properly function identification of targets is essential. Targets are receptors or proteins that play an important part in the cancer cell growth and survival. There are two main approaches to identify potentially good targets. The first one is associated to the amount of proteins located in cancer cells. If a specific protein is found only in cancer cells or in abundant quantities in them compared to healthy cells, then this protein could be a candidate target. An example of such a target is the human epidermal growth factor receptor 2 protein (HER-2) (Wikipedia, 2018). This factor is expressed at high levels on the surface of some cancer cells. Several targeted therapies are directed against HER-2, including trastuzumab (Pratt et al., 2015), which is approved to treat certain breast and stomach cancers that overexpress HER-2.

The second approach to identify possible targets is related to the ability of cancer cells to produce mutant proteins that are responsible for cancer progression. For example, the cell growth signaling protein BRAF is present in an altered form (known as BRAF V600E) in many melanomas (Cheng et al., 2018). Vemurafenib targets this mutant form of the BRAF protein and is approved to treat patients with inoperable or metastatic melanoma that contains this altered BRAF protein (Luke and Hodi, 2012). Chromosome abnormalities in cancer cells, also, are responsible for fusion gene creation. A fusion gene is when a gene incorporates parts of two different genes and their product the fusion protein might drive cancer development. These proteins can be targets of MTAs. An example is imatinib mesylate that targets the BCR-ABL fusion protein (Sharma et al., 2010), which is made from pieces of two genes that get joined together in some leukemia cells and promotes the growth of leukemic cells.

After identifying a possible target, a new therapy is developed aiming for this target. This therapy can reduce the activity of the target or prevent the target from binding to a receptor that it normally activates. That way it interferes with its ability to promote cancer cell growth or survival. Most targeted therapies are either small molecules or monoclonal antibodies. Small-molecule compounds are mainly developed for targets that are located inside the cell, whereas monoclonal antibodies are used for targets outside or on the surface of the cancer cell, due to their large size.

An issue with targeted therapies, also present in cytotoxic therapies, is the potential cancer cell resistance to the treatment. In other words, after a period of time the target might mutate and hence, the therapy can no longer interact with it. Another type of resistance is when the tumor finds an alternative pathway to achieve tumor growth that does not depend on the target. It is for that reason that targeted therapies are usually administered in combination either with another targeted therapy or with classic cytotoxic agents. Another issue with regard to MTAs is patient eligibility. For certain types of cancer, most patients with that cancer will have an appropriate target for a particular targeted therapy and, thus, will be candidates to be treated with that therapy. Nonetheless, there are cancers for which not all patients can profit from targeted therapies. Patients are eligible for a targeted therapy if for instance, their tumor has a specific gene mutation that codes for the target. If no such a mutation is present then the therapy will have nothing to target.

For targeted therapies, biomarkers play an important part. Their first role is to identify eligible patients, indicating who would likely respond to a treatment. An example of such a biomarker is HER2 for breast cancer. Patients with HER2/neu amplified tumors are more likely to respond to Trastuzumab, which is a monoclonal antibody. Another aspect of biomarkers is to measure patients' response to a specific treatment. When a good biomarker has been identified, we use this biomarker to measure the activity of the treatment. In oncology, we classically work on two types of biomarkers with different patterns of trajectories. (i) Biomarkers that are produced by the tumor, such as cancer antigen 125 (CA 125) in ovarian cancer, prostate specific antigen (PSA) in prostate cancer, and angiotensin-converting enzyme (ACE) in colorectal cancers, and reflect the tumor progression; they may decrease in case of response to treatment before re-increasing when the tumor escapes the treatment and progresses (Figure 1.4 a)). (ii) Other biomarkers reflect the direct action of the treatment such as the plasma concentration or the level of antibodies binding to their target. The trajectory would then be an increase in time followed by a decrease when the tumor gets resistant and subsequently progresses (Figure 1.4 b)). Most often, in oncology, we measure activity with tumor reduction alone or tumor reduction in combination with a biomarker. Even though tumor reduction is

a precise measure of treatment activity, it is evaluated with CT scans, ultrasounds or MRIs (magnetic resonance imaging), therefore, it is not possible to measure it at regular intervals (James et al., 1999).

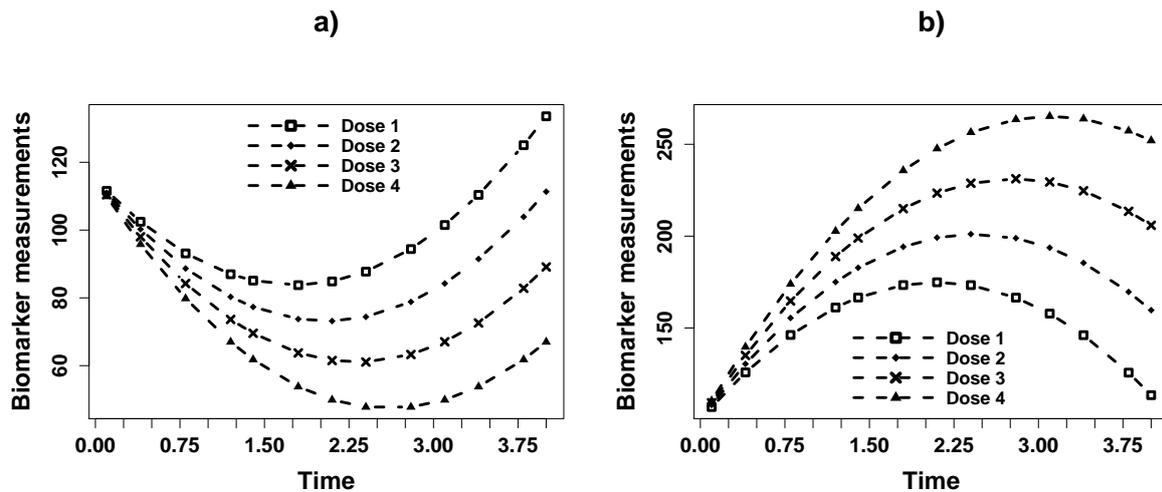


Figure 1.4: Types of biomarker trajectories.

1.5.2 Designs combining toxicity and activity

As mentioned before, targeted therapies require different techniques for dose finding, due to their biological activity and their administration route. In an attempt to address some of the concerns described in 1.1, new designs have been developed over the past few years, that consider as a primary endpoint both toxicity and activity (O'Quigley et al., 2001), targeting now the OD or the optimal biological dose (OBD), that may not necessarily be the MTD. Usually, the OD is defined as the dose that achieves high activity, as that defined by clinicians, while at the same time satisfying certain toxicity

requirements. The value of these designs is that they combine phase I and phase II designs and they make it feasible to recommend a dose, the OD, that might be as active as the MTD, but at the same time less toxic. In the next sections, we divide dose-finding designs in three large categories: (1) designs that consider toxicity and activity as binary endpoints, measured over one treatment cycle, (2) designs that consider toxicity and activity as binary endpoints, including time in their models, and (3) designs that consider binary toxicity and continuous activity outcomes.

1.5.2.1 Binary toxicity and activity

The majority of the proposed designs usually considers outcomes of interest to be binary variables. Toxicity is measured based on the DLTs (1 for DLT and 0 for non-DLT). Similarly, activity is a variable that is dichotomized, indicating patients who respond and do not respond to a treatment. For example, patients with a 30% tumor reduction are considered as responders, whereas anything below 30% is considered as a non-response.

In one of the landmark papers, Thall and Cook (2004) described an adaptive model based Bayesian procedure with a dose escalation scheme based on a “trade-off” between toxicity and activity. In particular, they allowed to escalate to doses that might be more toxic than the usual TTL as long as they produced a high probability of activity, as that defined by clinicians. Thall and Cook (2004) formulated their models in terms of marginal probabilities. For dose-toxicity relationship they assumed a monotonic increase

$$\pi_T(d, \boldsymbol{\theta}) = g^{-1}\{\eta_T(d, \boldsymbol{\theta})\},$$

with $\boldsymbol{\theta}$ being the model parameter vector and

$$\eta_T(d, \boldsymbol{\theta}) = a_0 + a_1d,$$

where a_0 and a_1 are unknown parameters. For the dose-activity relationship they allowed for more flexibility and non-monotone relations, by adding a quadratic term for dose.

The marginal probability is described as

$$\pi_A(d, \boldsymbol{\theta}) = g^{-1}\{\eta_E(d, \boldsymbol{\theta})\},$$

where

$$\eta_E(d, \boldsymbol{\theta}) = \beta_0 + \beta_1 d + \beta_2 d^2,$$

where β_0 , β_1 and β_2 are unknown parameters. They jointly modeled the two outcomes using a Gumbel copula model

$$\begin{aligned} \pi_{\lambda, \nu} &= (\pi_A)^\lambda (1 - \pi_A)^{1-\lambda} (\pi_T)^\nu (1 - \pi_T)^{1-\nu} \\ &+ (-1)^{\lambda+\nu} (\pi_A)(1 - \pi_A)(\pi_T)(1 - \pi_T) \left(\frac{e^\gamma - 1}{e^\gamma + 1} \right), \end{aligned}$$

where γ is the association parameter of the two models and $(\lambda, \nu) \in \{0, 1\}$. Then the likelihood function for a single individual is given by

$$\mathcal{L}(Y_A, Z_T, d | \boldsymbol{\theta}) = \prod_{\lambda=0}^1 \prod_{\nu=0}^1 \{\pi_{\lambda, \nu}(d, \boldsymbol{\theta})\}^{I((Y_A, Z_T) = (\lambda, \nu))},$$

where Y_A and Z_T are the indicators of activity and toxicity, respectively.

For the identification of the OD, Thall and Cook (2004) developed a “trade-off” algorithm that selects the dose with the best compromise between the probability of toxicity and of activity. Before the trial initiation they construct a target activity–toxicity “trade-off” contour, by fitting a curve to target values elicited from clinicians. Doses on the contour are considered to be “equally” desirable. The target contour is then used to estimate the desirability of each dose that is associated to some probability of toxicity and some probability of activity. For a better understanding we illustrate an example. To construct the contour clinicians elicit three target values $\{\pi_1^*, \pi_2^*, \pi_3^*\}$ that are considered

to be equally desirable. These target values are pairs of probabilities of activity and toxicity. The first target $\pi_1^* = (\pi_{1,A}, 0)$ is the acceptable activity in the absence of toxicity. The second target $\pi_2^* = (1, \max(\pi_{1,T}))$ refers to the maximum toxicity that clinicians are willing to accept if we have the maximum activity. Finally, the third target π_3^* is between the two first extreme targets. For our example (Figure 1.5) $\pi_1^* = (0.15, 0)$, $\pi_2^* = (1, 0.60)$ and $\pi_3^* = (0.25, 0.30)$.

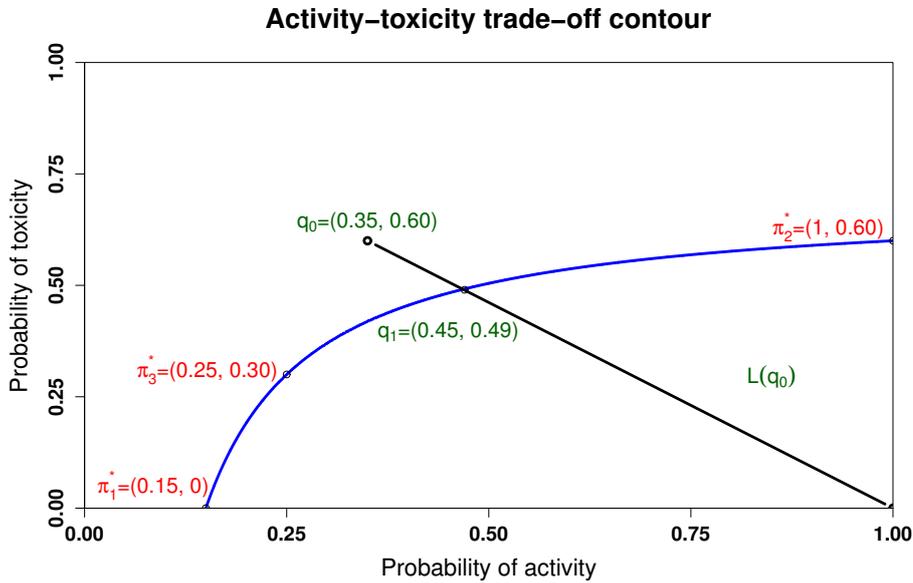


Figure 1.5: Activity-toxicity trade-off contour. The three elicited target points that determine the contour are given by π_1^* , π_2^* , and π_3^* . $L(q_0)$ indicates the line from q_0 to $(1, 0)$.

To design the contour any function can be used $\pi_T = f(\pi_A)$ that increases continuously over $\pi_{1,A}^* \leq \pi_A \leq 1$. For Figure 1.5 they used

$$f(x) = \delta_0 + \frac{\delta_1}{x} + \frac{\delta_2}{x^2}.$$

After designing the contour we use it to calculate the desirability of each dose level. Imag-

ine that we are interested in a dose with a pair of probabilities equal to $q_0 = (0.35, 0.60)$. First, we draw a line $L(q_0)$ from q_0 to $(1, 0)$ and find the point q_1 where $L(q)$ and the contour intersect. Then, we calculate the Euclidean distance $\rho(q_0)$ from q_0 to $(1, 0)$ and $\rho(q_1)$ from q_1 to $(1, 0)$. The desirability of the dose is calculated by

$$des(q_0) = \frac{\rho(q_1)}{\rho(q_0)} - 1.$$

A dose with a probability pair q between the contour and $(1, 0)$ will have a $des(q) > 0$, whereas outside of the contour region $des(q) < 0$. The dose that maximizes the desirability function is the dose administered to the next patient.

As previously discussed, the activity of MTAs increases monotonically with dose and contrary to cytotoxic therapies, it may reach a plateau after a certain dose level. This is of importance when modeling activity, since usually there is not merit in further escalating to doses located after the beginning of the plateau. Riviere et al. (2016) considered such a dose-activity relationship and introduced a Bayesian design that aimed to identify the OBD. The OBD is the dose with the lowest toxicity among those with the highest activity. Imagine that there are M planned doses and $D = \{d_{(1)}, \dots, d_{(M)}\}$. For the toxicity outcome they used a logistic model

$$\text{logit}(\pi_{(l)}) = a_0 + a_1 d_{(l)},$$

where $\pi_{(l)}$ is the probability of DLT at the l^{th} dose and a_0 and a_1 are unknown parameters. Let z_i denote the DLT outcome, (1 DLT and 0 no DLT), of the i^{th} individual, $i = 1, \dots, N$, treated at dose level l . Then, the likelihood after i patients have been enrolled is

$$\mathcal{L}(\mathcal{D}_T | a_0, a_1) = \prod_{m=1}^i \pi_{m(l)}^{z_m} (1 - \pi_{m(l)})^{1-z_m},$$

where $\mathcal{D}_T = \{(d_1, z_1), \dots, (d_i, z_i)\}$. For the activity outcome, they considered a dose-activity relationship that reaches a plateau after a certain dose level. Thus, they employed

a logistic model with a plateau parameter that can capture the increasing-then-plateau pattern of the dose-activity relationship of MTAs. Let $\phi_{(l)}$ denote the probability of activity at dose level l . Then,

$$\text{logit}(\phi_{(l)}) = \beta_0 + \beta_1(d_{(l)}\mathbf{1}_{(l < pl)} + d_{(pl)}\mathbf{1}_{(l \geq pl)}),$$

where β_0 and β_1 are unknown parameters, $\mathbf{1}$ is the indicator function, and pl is an index of plateau taking values between 1 and M . For all dose levels below pl we make the assumption that the dose-activity relationship increases monotonically, whereas for all doses equal or above pl the dose-activity relationship admits a plateau. Let y_i denote the activity outcome, (1 activity and 0 no activity), and $\mathcal{D}_A = \{(d_1, y_1), \dots, (d_i, y_i)\}$. The likelihood for the activity outcome is formulated as

$$\mathcal{L}(\mathcal{D}_A | \beta_0, \beta_1, pl) = \prod_{m=1}^i \phi_{m(l)}^{y_m} (1 - \phi_{m(l)})^{1-y_m}. \quad (1.3)$$

The authors accommodated for the case of delayed activity outcomes by weighting the likelihood 1.3 with the patient's follow-up time. Therefore, the likelihood for the activity outcome is

$$\mathcal{L}(\mathcal{D}_A | \beta_0, \beta_1, pl) = \prod_{m=1}^i (w_{m,i} \phi_{m(l)})^{y_{m,i}} (1 - w_{m,i} \phi_{m(l)})^{1-y_{m,i}}.$$

Riviere et al. (2016) selected the same weighting scheme as Cheung and Chappell (2000) for the TITE-CRM.

For the dose finding algorithm, Riviere et al. (2016) started with a dose escalation scheme, in which patients enter the trial in cohorts of three. The first cohort is treated at the first dose level and if no DLT is observed then a second cohort is allocated to the second dose level. They continue with this rule based design until they encounter the first DLT. Then they switch to the model based dose finding phase. This dose escalation scheme serves to collect some information before applying the Bayesian inference. For the

model based part consider c to be the cohort size and $n_{(l),i}$ to be the number of patients treated at dose level l . The dose level l is considered to be admissible if it satisfies a safety requirement

$$P(\pi_{(l)} > \zeta) < L_T,$$

and an activity requirement

$$P(\phi_{(l)} > \xi) \geq L_A 1_{(n_{(l),i} > \max(c,3))},$$

where ζ and ξ denote the pre-specified toxicity upper bound and activity lower bound, respectively. L_T and L_A are the respective probability thresholds for toxicity and efficacy. If we assume that \mathcal{B} contains all the admissible doses, then the $(i+1)^{th}$ will be allocated to the dose level with the highest activity in \mathcal{B}

$$d = \min \left(\operatorname{argmax}_{(l) \in \mathcal{B}} (\hat{\phi}_{(l)}) \right).$$

If several dose levels have the same activity, it means that activity reached a plateau and among them the dose with the lowest toxicity is selected. Even under the Bayesian framework, we see that it is preferred to start with a rule-based dose escalation scheme, so as to acquire some data before applying the Bayesian inference.

Below, we briefly present similar designs that can be found in the literature. Hunsberger et al. (2005) developed two types of designs for binary outcomes, with the goal to identify a biologically adequate dose. Toxicity was assumed to be relatively low and thus dose escalation was established only upon activity measurements, mimicking the ‘3+3’ design, till the moment a DLT occurred. A different approach that can be found in the current bibliography is the TriCRM (Zhang et al., 2006). This design considered a trinomial ordinal outcome of no activity and no toxicity, activity without toxicity and just toxicity, with these three probabilities summing up to one. It was assumed that the probability of no response at all decreased monotonically with dose, the probability of

toxicity increased monotonically with dose, whereas for the probability of response there could be a non-monotone relationship with dose. The dose-escalation process initiated with a screening based on toxicity measurements and from a safe range of doses the OBD was selected upon some predefined activity criteria. This approach allowed for early stopping due to safety constraints. Ivanova and Xiao (2013) compared three designs, with the purpose to investigate which one performs better under the assumption of plateau shaped functions for activity data. These three designs were a group design imitating the ‘3+3’ design, the t-statistic (Ivanova and Kim, 2009) and the classic CRM (O’Quigley et al., 1990). They showed that the t-statistic performed better than the other two, due to the fact that it allocated many patients to the estimated target dose and selected the lowest dose on the plateau quite often. Cai et al. (2014) proposed a Bayesian design that accommodated for non-monotonic dose-activity relationships. In particular, for the toxicity endpoint they assumed a change-point model accounting for the fact that the dose-toxicity surface of the combinational agents may plateau at higher dose levels. A flexible logistic model was implemented to capture the possible non-monotonic pattern of the dose-activity relationship. Three different models for the activity measurements were also compared by Zang et al. (2015). In particular they introduced a Bayesian based technique for a binary beta-binomial model for toxicity and three separate models for activity; a logistic (parametric), an isotonic (non parametric) and local logistic (semi-parametric). The isotonic and local logistic seemed to have good operating characteristics, but the main limitation of all three of them was that they could not address cases of delayed outcomes. Another design under the Bayesian framework that combined features of the CRM and the order restricted inference (Robertson et al., 1988) was introduced by Wages and Tait (2015). Toxicity was represented by a bivariate binary model, whereas for activity they made use of a class of working models. Specifically, they considered both unimodal and plateau skeleton models. Safety and futility rules for early stopping of the trial were also taken into consideration in the study design. A common goal among several designs was to reflect in their models the plateau shape of the dose-activity relationship.

1.5.2.2 Time-to-event outcomes

Even though the above designs included both toxicity and activity outcomes, only few of them included time in their models. Yuan and Yin (2009) developed a Bayesian adaptive design by jointly modeling time-to-DLT and time-to-activity, using a copula model. Their primary objective was to identify the OD, tackling the issue of delayed outcomes, so as to allow subjects to be enrolled in the trial, even when only partial information had been acquired from the previous cohorts.

Let t be the time from the initial treatment until the occurrence of a DLT. The hazard function for the time-to-DLT, under the Cox proportional hazards model is

$$h_T(t|d) = h_{0T}(t) \exp(ad),$$

where $h_{0T}(t)$ is the baseline hazard, modelled by a Weibull distribution, and a is the parameter for the dose variable. Then the survival function is given by

$$S_T(t|d) = \exp\{-\lambda_T t^{(\mu T)} \exp(ad)\},$$

where λ_T and μT are the scale and shape parameters of the Weibull distribution, respectively. The survival function for the time-to-activity outcome is

$$S_A(t|d) = \exp\{-\lambda_A t^{(\mu A)} \exp(\beta d)\},$$

where β is the parameter for the dose variable. Yuan and Yin (2009) jointly modeled toxicity and activity using a Clayton copula model, such that

$$S(t_T, t_A|d) = \{S_T(t_T|d)^{-1/\gamma} + S_A(t_A|d)^{-1/\gamma} - 1\}^{-\gamma},$$

where t_T and t_A denote time to toxicity and activity respectively, and $\gamma > 0$ measures the correlation of the two outcomes.

For the dose finding algorithm the authors suggested selecting the OD using the areas under survival curves (AUSC), see Figure 1.6. More precisely, they took the ratio of the AUSC of toxicity A_T and activity A_A up to a prespecified follow-up time τ ,

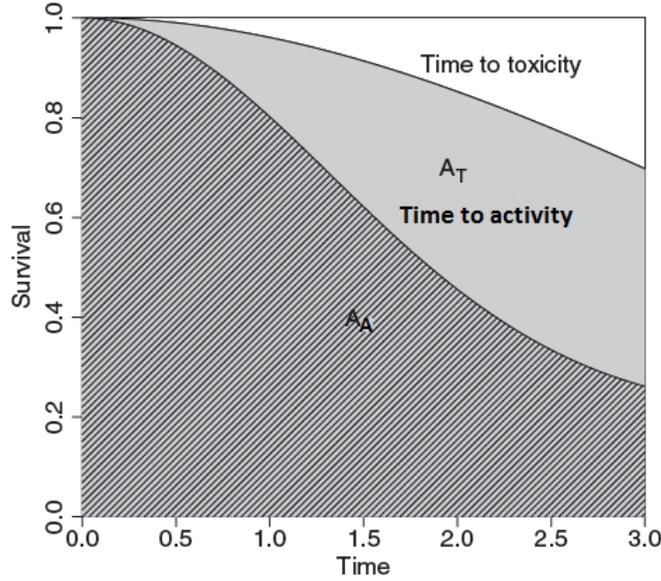


Figure 1.6: Illustration of the AUSCs corresponding to toxicity and activity.

$$\frac{A_T}{A_A} = \frac{\mu_T^{-1} \{\lambda_T \exp(ad)\}^{-1/\mu_T} \Gamma\{\mu_T^{-1}, \lambda_T \exp(ad)\tau^{\mu_T}\}}{(1 - \pi)\tau + \pi \mu_A^{-1} \{\lambda_A \exp(\beta d)\}^{-1/\mu_A} \Gamma\{\mu_A^{-1}, \lambda_A \exp(\beta d)\tau^{\mu_A}\}},$$

where Γ is the Gamma distribution and π is the fraction of patients that benefit from the treatment. Then, the OD is the dose that maximizes $\frac{A_T}{A_A}$.

An advantage of the AUSC ratio for the dose finding algorithm is that it takes into account the toxicity and activity rates, as well as how quickly patients experience toxicity or achieve activity. However, a limitation of their design is related to the start up phase of the design. To acquire some information, before applying the Bayesian inference, the

authors suggested initiating the trial with a rule based design that mimics the ‘3+3’ design. The related issue is that before a new cohort enters the study, the previous one should have been completely followed. We can imagine that this constraint can be a hindrance in cases where we expect to have a long follow-up period.

Similarly, Koopmeiners and Modiano (2014) modeled both endpoints as time-to-event outcomes and based their dose escalation decision rules on the activity-toxicity “trade-off” (Thall and Cook, 2004). Another design for delayed outcomes (Jin et al., 2014) treated late-onset responses as missing data. More specifically, Jin et al. (2014) proposed a design that handles the problem of delayed outcomes, by accounting for each patient’s follow-up time prior to evaluation of toxicity and activity outcomes, and treating outcomes that have not yet been observed as missing values. They used data augmentation to impute each missing outcome using partial follow-up times and complete outcome data. For the design’s decision rules, they combined observed and imputed data for all patients who have been treated.

1.5.2.3 Binary toxicity and continuous activity

Nonetheless, binary variables are known to be associated with irreducible variances and are often less informative than continuous variables. Nowadays, more and more biomarkers are used as indicators of activity in oncology (CA 125, PSA, circulating tumor cells, ACE, etc.). An approach of jointly modeling, under the Bayesian framework, binary toxicity and continuous activity outcome, measured at a unique timepoint, via a latent Gaussian variable, was originally suggested by Bekele and Shen (2005). An essential aspect of their article focused on the importance of modeling continuous and not binary biomarker measurements, indicating that dichotomization of the response variable can lead to substantial loss of information and as a corollary to inappropriate dose selection.

To describe the dose-toxicity relationship, Bekele and Shen (2005) employed a probit model. To allow for model flexibility, they assumed that there is a separate parameter for each dose level. They introduced a latent variable W_i for the i^{th} patient, which follows a normal distribution with mean $\mathbf{a}_w \mathbf{d}_i$ and variance 1, where $\mathbf{a}_w = (a_{w,(1)}, \dots, a_{w,(M)})^T$ is a $M \times 1$ vector of unknown parameters, M being the total number of doses. The latent

variable is related to toxicity Z_i via the condition

$$Z_i = \begin{cases} 0 & \text{if } W_i \in A_0 \\ 1 & \text{if } W_i \in A_1, \end{cases}$$

where $A_0 = (-\infty, 0]$ and $A_1 = (0, \infty)$. Then, the marginal probability of toxicity is

$$Pr(Z_i = 1 | d_{i(l)}) = \Phi(a_{W,(l)}),$$

where a_W is an unknown parameter. For the activity endpoint, we have a state-space model. Let Y denote the continuous biomarker variable and $\boldsymbol{\beta}_Y = (\beta_{Y,(1)}, \dots, \beta_{Y,(M)})^T$ define an unknown vector of parameters, where $\beta_{Y,(l)}$ reflects the mean biologic activity for the l^{th} dose. The state-space model assumes a recursive relationship between $\beta_{Y,(l)}$ and $\beta_{Y,(l-1)}$. Then the relationship between the two parameters is given by

$$\beta_{Y,(l)} = \beta_{Y,(l-1)} + u_{(l)}, \quad u_{(l)} \sim N(0, \sigma^2),$$

This model was motivated so as to consider a flexible dose-activity relationship, which does not necessarily increase monotonically. The authors joint modeled the two outcomes via a latent Gaussian variable. The joint distribution of (Y_i, Z_i) is assumed to have a bivariate normal distribution with mean $(\mathbf{d}_i \boldsymbol{\beta}_Y, \mathbf{d}_i \mathbf{a}_W)^T$ and variance covariance matrix Ω , such that

$$\Omega = \begin{pmatrix} \sigma_Y^2 + \rho & \rho \\ \rho & 1 \end{pmatrix},$$

that is positive definite and ρ is the covariance parameter between Y and Z , which induces the dependence between toxicity and activity. Given the observed data the likelihood for

the i^{th} patient takes the form

$$f(y_i, z_i | \beta_Y, a_W, \Omega, d_i) = \prod_{m=0}^1 \left\{ \int_{A_m} \phi_{Y,W}(y_i, w_i; (\mathbf{d}_i \boldsymbol{\beta}_Y, \mathbf{d}_i \mathbf{a}_W)^T, \Omega) dw_i \right\}^{I(z_i=m)},$$

where $\phi_{Y,W}(y, w, \mu, \Sigma)$ is a bivariate normal density with mean μ and variance covariance matrix Σ .

In the beginning of the trial the first cohort of patients is assigned to the first dose level. The size of the cohort c can vary between 1 and 6 patients. A dose is considered to be adequately tried, i.e. there is enough information concerning toxicity and activity of that dose, if at any point $n_{(l)} \geq c$. In other words, a dose is considered to be adequately tried, if more than one cohort has been treated to that dose level. Before the trial, a clinician defines the minimally acceptable biomarker level Y_{min}^* and the maximum acceptable toxicity probability π_T . Finally, let η_1 and η_2 be the prespecified probability thresholds. Then a set of doses is considered to be acceptable if it satisfies two requirements. For activity

$$Pr(\beta_{(l)} > Y_{min}^* | data, n_{(l)} \geq c) > \eta_1$$

and for toxicity

$$Pr(\Phi(a_{W,(l)}) < \pi_T | data, n_{(l)} \geq c) > \eta_2.$$

Therefore, a dose is defined as acceptable if the posterior probability of the mean biomarker level is greater than Y_{min}^* is greater than η_1 and if the posterior probability of toxicity is less than π_T is greater than η_2 . If more than one dose is deemed acceptable then the dose administered to the next cohort of patients is selected based on a rule similar to that proposed by Thall and Cook (2004). Let $p_y(d_{(l)}) = E(\beta_{Y,(l)} | data)$ be the expected posterior mean for the biomarker and $p_z(d_{(l)}) = E(\Phi(a_{W,(l)}) | data)$ be the expected posterior probability of toxicity at dose $d_{(l)}$. Finally, we elicit from clinicians the largest possible

biomarker value Y_{max}^* and define $(Y_{max}^*, 0)$ to be the optimal point, which corresponds to the maximum activity and zero toxicity. Then, for each of the acceptable doses we estimate their Euclidean distance $g_{(l)}$ from the optimal point.

$$g_{(l)} = \sqrt{\left[\frac{Y_{max}^* - p_y(d_{(l)})}{Y_{max}^*}\right]^2 + p_z(d_{(l)})^2}.$$

A dose that minimizes the distance from the optimal point is considered to be the OD, allocated to the next cohort.

A similar design to that of Bekele and Shen (2005) was proposed by Hirakawa (2012). On the same page, Yeung et al. (2015) introduced an adaptive design for the identification of the OD, that combines toxicity and activity through a logistic regression and a linear log-log regression model, respectively. The activity model was restricted only to patients with no DLTs. Dose escalation was achieved via a gain function or otherwise a “trade-off” between the two outcomes. Stopping rules were applied either for safety or when the OD was reached. Finally, Yeung et al. (2017) proposed a similar design with the one described above, only in this case activity was modeled via a flexible non-parametric model.

1.5.3 Future concerns

Despite the growing popularity of designs that model both toxicity and activity outcomes, there are still important aspects of phase I trials that need to be considered. As mentioned before, the prolonged administration of MTAs can lead to cumulative toxicities. Ignoring this information when designing a trial may result in selecting a dose that could eventually be highly toxic in the long-term. Unfortunately, given the current designs that can only be investigated in phase II or III trials, in which follow-up is significantly longer. It is vital to define the MTD or the RP2D based on DLTs occurring after multiple cycles and not just the first one. Furthermore, most of the methods consider time as a continuous variable. However, in most clinical trials the exact time of the event

is not known. Severe toxicities are recorded at the time of a patient's visit, which is usually at the end of the treatment cycle or the beginning of the next one. As a result, toxicity data are censored by intervals, which most often are the treatment cycles. If we account for this censored information it is feasible to measure the cumulative probability of toxicity of each dose at each treatment cycle. An additional point here is with respect to the activity measurements. When a certain biomarker is used as an indicator of activity then for each patient we have repeated measurements during the follow-up. These measurements are collected in regular time intervals in order to explore the biomarker trajectory till disease progression or excessive toxicity, both of which would force a patient to drop-out of the study. All methods described here consider the activity at a unique time point. One major limitation of this approach is that we cannot know in advance which measurement is the most informative, so selection is usually random. Secondly, by ignoring the time-activity relationship of the biomarker we discard a vital part of the information.

Chapter 2

An adaptive design for the identification of the optimal dose using joint modeling of continuous repeated biomarker measurements and time-to-toxicity

2.1 Introduction

In Chapter 1 we introduced the existing methods for dose-finding in oncology. Discussion also revolved around challenges that we face when designing phase I trials, associated with cumulative toxicities, repeated measurements of biomarkers of activity as well as missing data. In this chapter, we attempt to address these important challenges by introducing a new adaptive design.

The remainder of Chapter 2 is organized as follows. In Section 2.2 we present the objective of our work along with a motivating example from a clinical trial (Subsection 2.2.1). Section 2.3 is the methods part, where we introduce the models for toxicity and activity and we describe the joint modeling framework. We talk about estimation difficulties, the study design and finally, the decision process for the identification of the

OD. Section 2.4 is dedicated to the simulation study we conducted in order to test the robustness and the operating characteristics of our design under various settings. We analytically present the simulation framework, i.e. how simulation scenarios were developed using the motivating example. Section 2.5 refers to the simulation results, including tables and also the sensitivity analysis. Discussion follows in Section 2.6.

2.2 Objective

We propose an adaptive design for the identification of the OD, in phase I/II trials of MTAs. The OD is defined as the lowest dose, within a range of active doses, that is below or equal to the MTD. The MTD is the dose with a cumulative risk of toxicity, or more precisely cumulative risk of DLT, over six cycles that minimizes the distance from a predefined TTL. The TTL is also defined cumulatively over six cycles. We employ a recent joint modeling technique of longitudinal continuous biomarker measurements and time-to-DLT data, using a shared random effect and likelihood approach. The biomarker is measured on a continuous time scale, whereas DLTs are measured on a discrete time scale.

2.2.1 Motivating example

This work is motivated by a french multicenter phase IB trial in patients with platinum resistant epithelial ovarian carcinoma. The aim of this trial is to assess the performance of the combination of three different agents, immunotherapy, biotherapy, and chemotherapy. The primary objective of the study is to determine the MTD and the OD that will be recommended for phase II. Activity of the drug combination is mainly evaluated with tumor volume and CA 125 measurements. CA 125 is a biomarker routinely used to assess response and progression in ovarian cancer clinical trials. According to the Gynecological Cancer Intergroup guidelines (Rustin et al., 2011), response is defined on the basis of a

reduction in the levels of CA 125. Progression is indicated by progressive serial elevation of serum CA 125, either above 70 units/ml for patients with normalization of CA 125 or otherwise above twice the nadir value. Zhou et al. (2016) provide an example of the CA 125 trajectory from a randomized clinical trial on ovarian cancer patients (Figure 2.1). We have simulated a similar example (Figure 2.2), using a linear mixed effects model, with a quadratic term for time that we shall use in the following. A maximum of eight cycles has been planned, with a duration of three weeks each, and a maximum of 42 patients is expected to be included in the study. Once the MTD is determined, then, based on both toxicity and activity data the OD will be defined. The OD will either be the same as the MTD or below the MTD, but will not exceed the MTD.

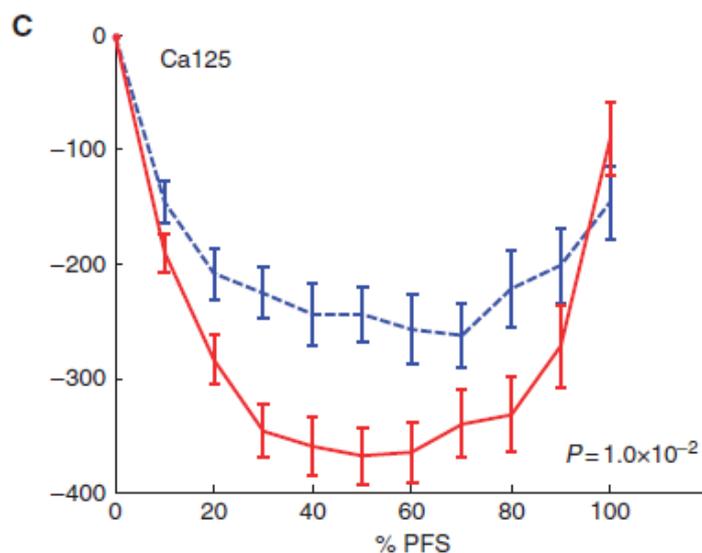


Figure 2.1: CA 125 trajectory plotted against the percentage of progression free survival. Red line refers to the experimental arm and blue line refers to the standard arm. Figure from a randomized clinical trial (Zhou et al., 2016).

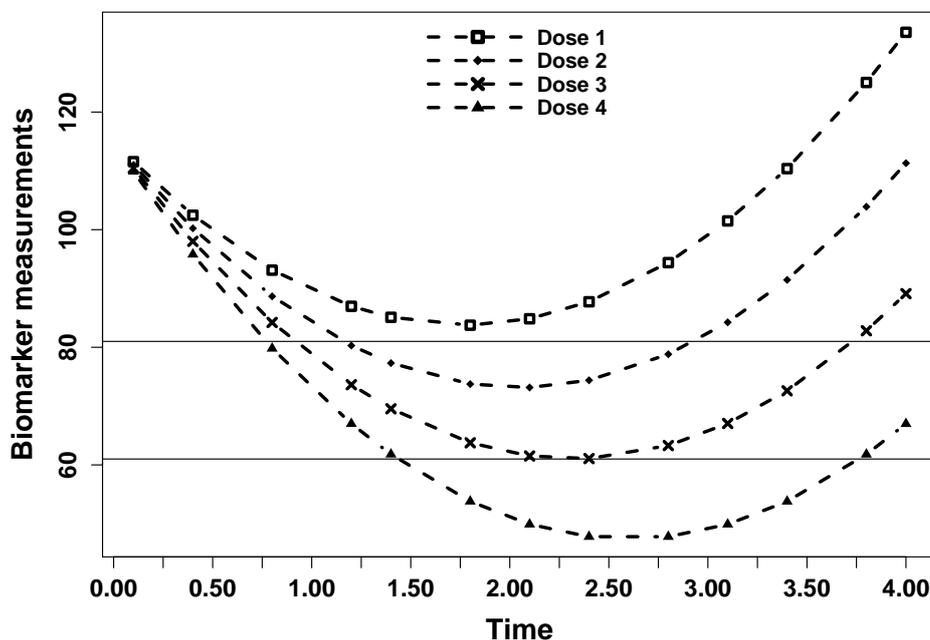


Figure 2.2: Simulated example of CA 125 trajectory. Horizontal lines represent the equivalence range between the MTD and the optimal dose.

2.3 Methods

2.3.1 General framework

Consider a trial of N individuals, $i = 1, \dots, N$ and a total of M planned doses to be tested $D = \{d_{(1)}, \dots, d_{(M)}\}$. Patients are sequentially assigned to a specific dose level, selected based on the data from previous patients. Cycles of treatment are repeatedly administered until DLT, disease progression, defined by an elevation of the biomarker, or consent withdrawal. At the end of each cycle study participants are evaluated as to whether they had a DLT during the cycle or not. We then have discrete time failure data. During the trial, the biomarker levels are measured at each weekly visit, in order to

monitor the treatment effect. As a result, for each patient we have repeated continuous biomarker measurements. One or more doses are considered to be “equally” active to the MTD if their predicted activity falls within a predefined equivalence range. For example, Figure 2.2 illustrates such an equivalence range. The minimum biomarker value, reflecting the maximum activity, of doses 2 and 3 fall within this equivalence range, indicating that the two doses are “equally” active. The OD is the lowest dose on the range of doses that are “equally” active to the MTD. If a DLT is observed or the agent is not active, based on the biomarker data, then the patient is removed from the study, resulting in censored DLT data and missing at random data (MAR) for the biomarker.

2.3.2 Modeling time-to-DLT

The time-to-DLT is a survival endpoint. Specifically, we are interested in the time of the first DLT of each patient. As mentioned in Section 1.5.3 DLTs are usually recorded at the end of a cycle. Consequently, for the toxicity outcome we have a discrete time scale $C = \{1, \dots, k\}$, where k is the total number of treatment cycles. To model this outcome a probit time-to-DLT model, adjusted for cycle and dose effects, is assumed. Let S_i be the time-to-DLT, so that

$$P(S_i = s | S_i > s - 1) = 1 - \Phi(a_0 + a_1 c_{i(s-1)} + a_g d_i), \quad c_{i(s-1)} \in \mathcal{C}, \quad (2.1)$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution and a_0 and a_1 are unknown parameters and d_i is the patient’s dose. Cycle $c_{i(s-1)}$ is included in the model, assuming a linear relationship on the probit scale for the hazard of DLT and the cycle. No time effect is assumed in the first cycle, as the intercept captures the hazard at the first cycle. In the spirit of the CRM (O’Quigley et al., 1990), we fix the slope of the dose to a constant value a_g . Given the model, the probability of having a DLT at cycle s is conditional on having no DLT in all the past cycles.

As mentioned above, the MTD is defined based on the cumulative probability of

DLT. We consider a total of 6 cycles in the study. The probability of DLT at each cycle is obtained as follows:

$$P_{(l)}(S_i = s) = \begin{cases} (1 - \Phi_{(l)1}) & s = 1 \\ \Phi_{(l)1}\Phi_{(l)2} \dots \Phi_{(l),s-1}(1 - \Phi_{(l)s}) & s = 2, \dots, 6 \\ \Phi_{(l)1}\Phi_{(l)2} \dots \Phi_{(l),s-1}\Phi_{(l)s} & s > 6, \end{cases}$$

where $1 - \Phi_{(l)s}$ is the hazard function at cycle s and dose level l , l being the index for the dose. Then, the cumulative probability of DLT at the end of cycle 6 is

$$P_{(l)}(S_i \leq 6) = \sum_{s=1}^6 P_{(l)}(S_i = s).$$

2.3.3 Modeling activity

For the activity endpoint we suppose a continuous timescale, such that $\mathcal{T} = (0, k)$. Let Y_{ij} be the repeated continuous biomarker measurements for the i^{th} individual, at times $t_{ij} \in \mathcal{T}$ and $j = 1, \dots, n$ represents the visit number. The model structure is motivated by the activity trend over time, shown in Figure 2.2. We assume that doses of treatment modify the course of CA 125 as follows

$$y_{ij} = \beta_0 + \beta_1 t_{ij}^2 + \beta_2 t_{ij} d_i + \beta_3 t_{ij} \log(\sqrt{d_i} + 1) + U_i t_{ij} + r_{ij}, \quad (2.2)$$

where β_0 , β_1 , β_2 , and β_3 are unknown parameters to be estimated and d_i , as before, is the dose. U_i refers to the random effect associated with the time of the longitudinal measurements and $U_i \sim N(0, \sigma_1^2)$. The model residuals, R_{ij} , are mutually independent and follow a Gaussian distribution, $MVN(\mathbf{0}, \sigma_2^2 I)$, with I being the identity matrix. Within our model we do not include a main effect for the dose, since baseline activity is not associated with the dose level, to which a patient is allocated to. Furthermore, we have adjusted the model for two interaction terms. The first one, $t_{ij} d_i$ is used to capture

the dose-time interaction, usually present in a clinical trial, whereas the second one $t_{ij} \log(\sqrt{(d_i)} + 1)$ permits for a smoother dose-activity relationship. In other words, the log square root gives a curve that approximates a plateau in the dose-activity relationship. Regarding the biomarker trajectory over time, we selected the square of the time t_{ij}^2 , to model the parabolic time-activity relationship (Figure 2.1). This is no the first time that a model like 2.2 is used for a biomarker. For example, Rizopoulos (2012) modeled the level of serum bilirubin accounting for a quadratic time-trend, as well as for interactions with t and t^2 , in order to capture the biomarker evolution in time. Phase I patients might respond to the treatment but will eventually progress, due to advanced disease, and the biomarker will increase back to the initial values or even higher. This is also clear by the CA 125 guidelines that state that patients should be censored, due to lack of activity once their biomarker values increase above a certain threshold.

In principle, the timing of the individual measurements can differ per patient, but in our design we assume fixed time points t_{ij} . It is implicit that the number of measurements varies per subject due to censoring. The dose parametrization makes the model suitable for monotonic or plateau shaped dose-activity relationship, but not for a parabolic one.

In this design we make the assumption that patient discontinuation is solely related to the set of observed responses. That could refer to either disease progression or occurrence of a DLT. Thus, missing responses correspond to MAR. Longitudinal models, estimated with maximum likelihood estimation, provide valid inferences under MAR given the fact that they are correctly specified (Rizopoulos, 2012) .

2.3.4 Joint modeling

When implementing a survival model the main focus is to investigate how certain factors can affect the overall survival. These factors are usually constant during the follow-up and are measured at baseline. An example of such a variable is the dose administered to each patient, in model 2.1 that is not modified during the course of the trial. Furthermore, in oncology, we know that toxicity is associated to drug activity

(Abola et al., 2014). Patients who respond to the treatment are more likely to present with a DLT. Therefore, when estimating the probability of toxicity, it is of interest to take into account this association. However, activity is a time-dependent variable. Time-dependent is a covariate that changes in time and this change may have an impact on the outcome. There are two types of time-dependent variables; the exogenous and the endogenous. Exogenous is a variable that can be measured at any time s . An example of an exogenous variable is the air pollution or the seasonal patterns that are associated with allergies or asthma attacks. Irrespectively of the event of interest or the subject under study, at any time we can measure the levels of the air pollution. On the other hand an endogenous variable cannot be measured at any time s . The main reason is its dependence on the subjects under study. In this study CA 125 biomarker measurements is an endogenous variable. CA 125 is an indicator of activity and patients are removed from the study after CA 125 increase, suggesting disease progression. Biomarker measurements are directly connected to patients, therefore if the latter are out of study information is no longer available. Consequently, there is missing data in the time-to-DLT outcome due to a factor not directly related to toxicity. An endogenous variable requires the survival of the subject for its existence and adjusting for it is not enough to capture its effect. That being said it is clear that measuring the effect of these variables requires alternative methods.

When the primary focus of an analysis is to investigate the association of a survival outcome with time-dependent endogenous variables, an alternative framework has been proposed known as joint modeling. The main idea behind this approach is that the outcomes of interest are estimated using a joint distribution (Schluchter, 1992; Faucett and Thomas, 1996; Wulfsohn and Tsiatis, 1997; Henderson et al., 2000; Tsiatis and Davidian, 2004; Rizopoulos, 2012). An idea of joint models is given in Figure 2.3 (Rizopoulos, 2012). In the top panel we have a hazard function that describes how the risk of an event, here toxicity, changes in time. In the bottom panel we see the longitudinal process. The asterisks are the longitudinal measurements, here the biomarker measurements, and the green line delineates the fitted model. Joint models posit that the hazard function at any time point is associated with the value of the longitudinal process at the same time

point. Several joint models have been proposed in the current literature, such as the joint models with shared random effects or joint latent class models (Ibrahim et al., 2010; Mccrink et al., 2013; Proust-Lima et al., 2014; Lawrence Gould et al., 2015). Estimation of these models can be achieved under both the maximum likelihood or the Bayesian framework.

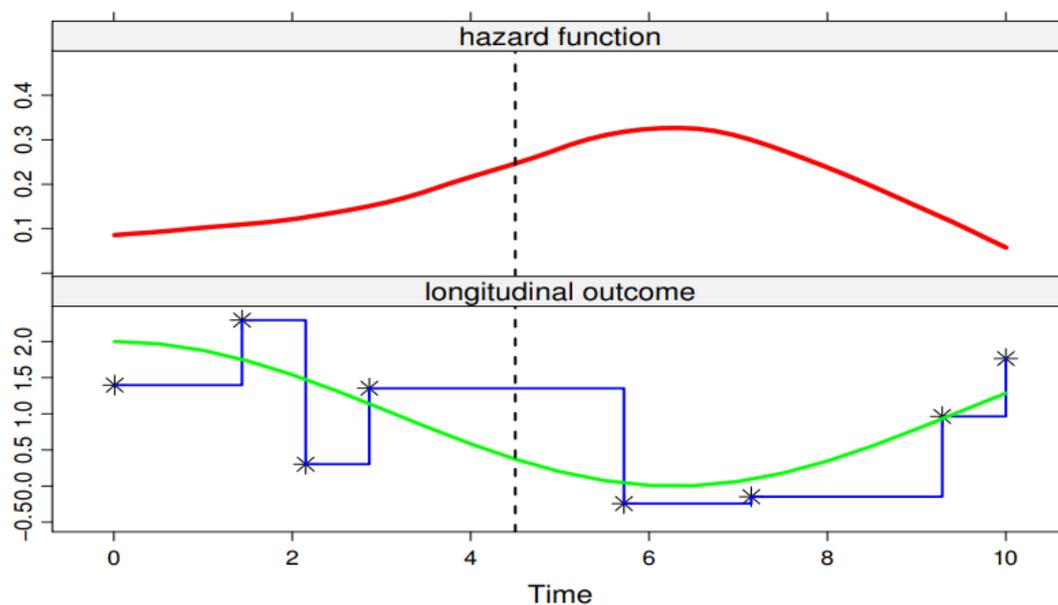


Figure 2.3: Top panel: Hazard function for the time-to-event outcome. Bottom panel: Longitudinal process (Rizopoulos, 2012).

Joint models have significantly gained ground over the past years because they benefit the analysis of both longitudinal and survival outcomes. To begin with, the use of the longitudinal model to incorporate the effect of the time-dependent variable on the survival outcome has been shown to provide less bias estimates than independent or separate model estimation (Ibrahim et al., 2010). Furthermore, integration of the survival infor-

mation into the evaluation of the longitudinal process directly embodies the effect of an informative missing-data mechanism. Joint models provide accurate inference, especially in the case where outcomes are strongly associated.

Returning to our framework, we shall focus on a joint modeling technique with a shared random effect, proposed by Barrett et al. (2015). This technique makes use of a probit time-to-DLT model 2.1 and a linear mixed effects model 2.2 that share the random effect that comes from the latter. Hence, model 2.1 takes the form

$$P(S_i = s | S_i > s - 1, U_i) = 1 - \Phi(a_0 + a_1 c_{i(s-1)} + a_g d_i + \gamma U_i), \quad c_{i(s-1)} \in \mathcal{C}, \quad (2.3)$$

where U_i , as before, is the random effect of the activity model and γ is the parameter associating the time-to-DLT outcome with the random effect. The underlying assumption of a shared random effect model is that repeated measurements and DLT times are independent conditionally on U_i .

A major challenge when jointly modeling time-to-event and longitudinal data is overcoming estimation difficulties, especially with small sample sizes. As underlined by Rizopoulos (2012), parameter estimation of shared random effect approaches for longitudinal and survival data is impeded by both numerical approximation of the log-likelihood integrals and numerical optimization. An attractive and elegant feature, of the joint modeling technique we use, is that it allows for likelihood inference. Barrett et al. (2015) chose the probit and the linear mixed effects model to take advantage of specific properties of the skew normal distribution (Azzalini, 2005; Arnold, 2009). Their aim was to express the joint likelihood in a closed form, by integrating out the random effects. They called it exact likelihood. Hence, the estimates of the model parameters are obtained directly from the numerical maximization of the likelihood.

Hereafter, for the notation we remain as close as possible to Barrett et al. (2015), thus, we omit the subscript i from the notation. Suppose now that there is an indicator δ , for the survival events, such that $\delta = 1$ for DLT and $\delta = 0$ for censoring. Define two vectors that summarize the model parameters, $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3)$ and $\boldsymbol{a} = (a_0, a_1, a_g)$ of the longitudinal and the survival model, respectively. Moreover, let X summarize covariates

that are operative on Y , i.e. time visits and dose and similarly let \tilde{X}_s be an $s \times 2$ matrix summarizing covariates up to time s , i.e. cycle and dose. Here, we make the assumption that survival variables are not time dependent. For simplicity purposes let $\boldsymbol{\theta}$ be a vector containing all unknown parameters and y be the realization of Y . After integrating out the random effects, the joint likelihood takes the form:

$$L(\boldsymbol{\theta}; y, s, \delta) = L_1(\boldsymbol{\theta}; y) L_2(\boldsymbol{\theta}; y, s, \delta). \quad (2.4)$$

The first term of the right part of expression (2.4), L_1 , corresponds to the longitudinal model and is expressed as

$$L_1(\boldsymbol{\theta}; y) = \frac{1}{(2\pi)^{n/2} \sigma_2^n |\sigma_1^2 H|^{1/2}} \exp\left\{-\frac{(y - X\boldsymbol{\beta})^T (y - X\boldsymbol{\beta})}{2\sigma_2^2} + \frac{h^T H h}{2}\right\},$$

where

$$H = \frac{t^T t}{\sigma_2^2} + (\sigma_1^2)^{-1}, \quad h = \frac{H^{-1} t^T (y - X\boldsymbol{\beta})}{\sigma_2^2}, \quad \text{and} \quad t = (t_1, \dots, t_n)^T.$$

The second term of expression (2.4), L_2 , is related to the event time model. Let Γ_s be an $s \times 1$ matrix that takes the value of γ and denotes the first s rows. Then L_2 for a single random effect can be written as

$$L_2(\boldsymbol{\theta}; y, s, \delta) = \begin{cases} \Phi^{(s)}\{\tilde{X}_s \mathbf{a} + \Gamma_s h; \mathbf{0}, I + \Gamma_s H^{-1} \Gamma_s^T\} & \text{for } \delta = 0, \\ \Phi^{(s-1)}\{\tilde{X}_{s-1} \mathbf{a} + \Gamma_{s-1} h; \mathbf{0}, I + \Gamma_{s-1} H^{-1} \Gamma_{s-1}^T\} \\ -\Phi^{(s)}\{\tilde{X}_s \mathbf{a} + \Gamma_s h; \mathbf{0}, I + \Gamma_s H^{-1} \Gamma_s^T\} & \text{for } \delta = 1, \end{cases}$$

where $\Phi^{(s)}$ is the s -dimensional cumulative normal distribution. The process to reach the analytical expression of the above relationships can be found in Barrett et al. (2015).

This method is highly robust in terms of parameter estimation and the authors demonstrate good mean square error and coverage. However, because of the small sample sizes

in phase I trials, we further examined its robustness under six different sample sizes, $N_1 = 15$, $N_2 = 20$, $N_3 = 25$, $N_4 = 30$, $N_5 = 40$, and $N_6 = 60$ and under the case of unbalanced data over time. We concluded to satisfactory bias and coverage. Results can be found in Table A.1 (Appendix A).

2.3.5 Study design

To maximize the likelihood, it should be non-monotone, so that the estimates do not arise on the boundary of the parameter space. In other words, we cannot proceed with the parameter estimation before some heterogeneity is encountered in the responses (O'Quigley and Shen, 1996). Even then, certain issues may arise when estimating the Hessian matrix for the standard errors. These difficulties are usually related to the small sample size. Even if optimization criteria are satisfied, the Hessian matrix may not be positive definite due to zero or negative eigenvalues. We then considered that fitting the joint model was possible only if certain criteria were satisfied; 1) at least 16 patients had been enrolled, 2) the standard errors of all joint model parameter estimates were below 20, and 3) the Hessian matrix was positive definite. Therefore, we split the design into two stages (Figure 2.4): for the first one dose escalation proceeds based on toxicity only and for the second one patient allocation is based on both outcomes and we identify the OD.

For the first stage, we introduce an escalation scheme that relies on a rule-based design, the '2+2' design. Patients enter the study in cohorts of two and dose escalation starts from the lowest dose level. In this stage interest lies on the toxicity observed in the first treatment cycle only. First, a cohort is assigned to dose level $d_{(l)}$. Then, at the end of a treatment cycle if no DLT has occurred the next cohort will be allocated to dose $d_{(l+1)}$ and if two DLTs have occurred then we de-escalate to dose $d_{(l-1)}$. Finally, if one DLT is observed then we remain at the same dose level.

When two DLTs and two non-DLTs have occurred, we stop the '2+2'. If all 3 criteria, mentioned above, are satisfied, we switch to the joint modeling stage. If not, then we rely

on a probit time-to-DLT CRM to guide dose-escalation. This approach, an extension of the CRM, is very similar to the one proposed by Sinclair and Whitehead (2014). Patients are sequentially enrolled in the trial at the best current estimate of the MTD, defined as the dose associated with some cumulative risk of DLT over six cycles. During this stage we accumulate the toxicity data from all treatment cycles and update the information before the entry of a new patient. Estimation relies on the same model 2.1 of the toxicity outcome

$$P(S = s | S > s - 1) = 1 - \Phi(a_0 + a_1 c_{(s-1)} + a_g d), \quad c_{(s-1)} \in \mathcal{C}. \quad (2.5)$$

Two DLTs and two non-DLTs, in two different cycles, are the minimum requirement for parameter estimation of model 2.5. We continue with the CRM until we can switch to the joint modeling. The ‘2+2’ and the CRM constitute the first stage of the design, where dose escalation depends solely on toxicity. Of note, any design could be applied in order to collect enough data to fit the joint model.

Finally, when all 3 criteria are met, we initiate the second stage and we update activity and toxicity estimates after each new patient, using joint modeling. Patients may enter into the study even if previous patients have not completed their six-cycle treatment period. Based on those estimates, we determine the MTD and the OD that is the dose the new patient will be allocated to. The OD allocated to the last patient in the trial is the dose recommended for phase II or III trials. For safety reasons dose skipping is not allowed. The trial ends either after all planned patients have been treated or due to excessive toxicity at the first dose level.

To avoid treating patients with a drug that would be excessively toxic at all dose levels, we implemented an early stopping rule, proposed by Ivanova et al. (2005). We generated a Pocock-type boundary for repeated testing of the toxicity, occurring at the first treatment cycle. A probability above 35% at the 10% level (one-sided) led to stop the trial. During the ‘2+2’ it is not possible to estimate the probability of DLT using model 2.3, thus we chose a non-parametric stopping rule that is applicable to both the ‘2+2’ and the CRM. We only considered DLTs at the first cycle of treatment for similar

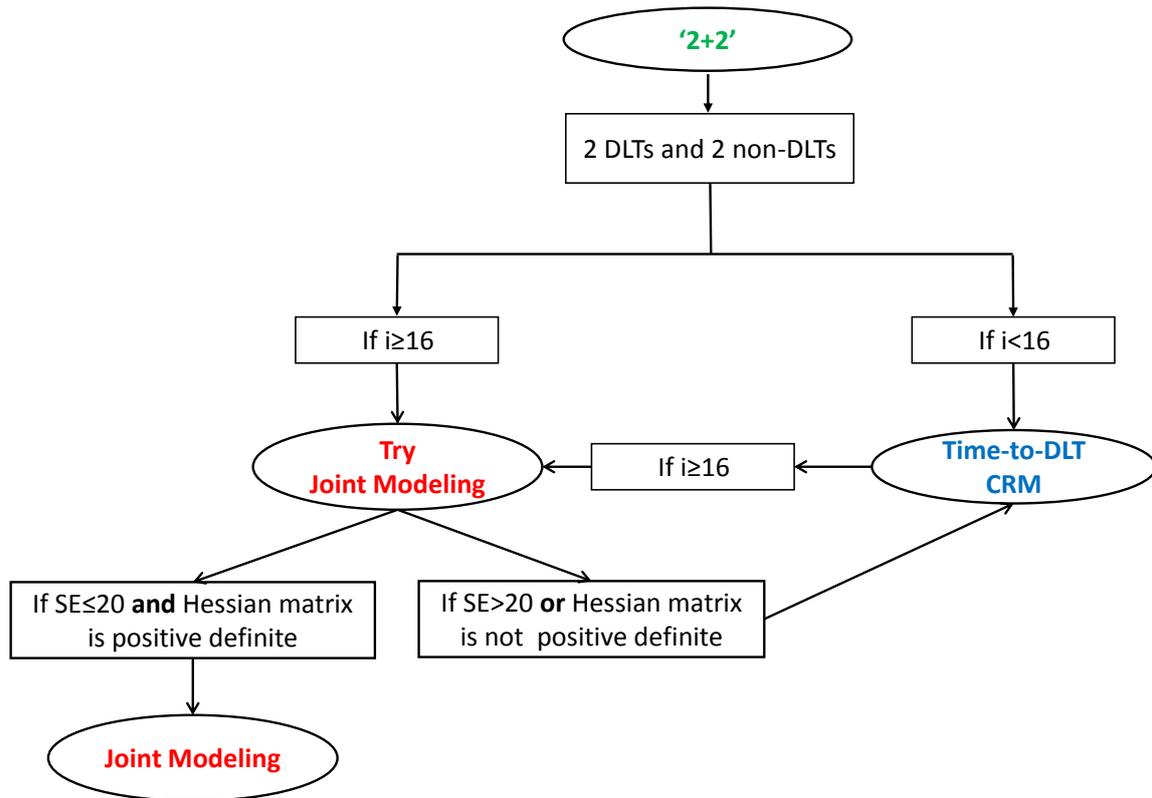


Figure 2.4: Representation of the study design described in Section 2.3.5. Abbreviations: DLTs, Dose Limiting Toxicities; i , index for the patient; SE, Standard Error; CRM, Continual Reassessment Method.

reasons. The explicit boundaries can be found in Table A.2 .

2.3.6 Decision process

Based on the joint model estimates, we determine the MTD and the OD. We start from the MTD. Taking the conditional probability of toxicity for $U = 0$ we minimize over

all dose levels

$$MTD = \min_{(l)} |\epsilon - P_{(l)}(S \leq 6|U = 0)|,$$

where ϵ is the TTL. Then the MTD is used to select a subset of safe doses B . This subset contains the MTD and all doses below the MTD, if the MTD is not the first dose level.

Then, from B we eventually select the OD, that is the lowest dose, from a range of doses, Λ , where $\Lambda \subseteq B$, that are “equally” active as the MTD. Particularly, for every $d_{(l)} \in B$ we estimate, at each time visit t_j and for each dose level l , the conditional predicted biomarker activity

$$\hat{y}_{j(l)} = \hat{\beta}_0 + \hat{\beta}_1 t_j^2 + \hat{\beta}_2 t_j d_{(l)} + \hat{\beta}_3 t_j \log(\sqrt{d_{(l)}} + 1).$$

The maximum drug activity is associated with the minimum biomarker measurement over time, thus, for each dose we select the minimum of the predicted values $\hat{y}_{min(l)}$. We identify doses that are “equally” active to the MTD as follows,

$$|\hat{y}_{min(MTD)} - \hat{y}_{min(l)}| \leq \zeta,$$

where ζ is the equivalence range. The lowest dose in Λ is selected for the next patient. Estimation of the MTD and the OD is repeated before the enrollment of a new patient.

For a better understanding, we summarize the above procedure by simulating an example. Consider a trial of $k = 4$ treatment cycles, one patient visit per cycle, so $n = 4$ and $M = 4$ planned doses. The estimations of the joint model parameters, serve to estimate the cumulative probability of DLT, over 4 cycles at each dose, as shown in Table 2.1. If we assume that the target is $\epsilon = 0.40$, then $MTD = d_{(3)}$, Figure 2.2. Consequently, the safe subset of doses is $B = \{d_{(1)}, d_{(2)}, d_{(3)}\}$. Next, for every $d_{(l)} \in B$ we select the minimum predicted value in time, $\hat{y}_{min(l)}$, which is depicted in bold in Table 2.2. We define $\zeta = 20$, that is any dose $d_{(l)} \in \Lambda$ to be “equally” active to the MTD, if $|\hat{y}_{min(MTD)} - \hat{y}_{min(l)}| \leq 20$. For instance, for dose level 2 we have $|\hat{y}_{min(MTD)} - \hat{y}_{min(2)}| = 13$,

meaning that $d_{(2)}$ and the MTD are “equally” active. On the contrary, for dose level 1 $|\hat{y}_{min(MTD)} - \hat{y}_{min(2)}| = 24$, thus, $d_{(1)}$ and the MTD are not “equally” active. Therefore, the OD is $d_{(2)}$. This can be seen in Figure 2.2, where the two horizontal lines delineate the range where doses are not clinically different from the MTD.

Table 2.1: Probability of DLT at each cycle $P(S = s)$ and cumulative probability of DLT over 4 cycles $P(S \leq 4)$ per dose level.

	$P(S = 1)$	$P(S = 2)$	$P(S = 3)$	$P(S = 4)$	$P(S \leq 4)$
$d_{(1)}$	0.05	0.03	0.01	0.01	0.10
$d_{(2)}$	0.12	0.07	0.04	0.02	0.25
$d_{(3)}$	0.22	0.11	0.08	0.04	0.45
$d_{(4)}$	0.30	0.20	0.14	0.06	0.70

Table 2.2: Predicted biomarker activity $\hat{y}_{min(l)}$ per cycle and per dose level. Bold values indicate the minimum predicted value over all cycles and for the safe doses only.

	$\hat{y}_{min(1)}$	$\hat{y}_{min(2)}$	$\hat{y}_{min(3)}$	$\hat{y}_{min(4)}$
$d_{(1)}$	102	85	88	110
$d_{(2)}$	100	77	74	91
$d_{(3)}$	98	70	61	73
$d_{(4)}$	96	62	48	54

2.3.7 Estimands

Joint model with shared random effect is used to obtain unbiased estimates of the expectation of the biomarker effect over time, while accounting for attrition due to severe toxicity and vice versa. The immortal cohort is not very different from the intent to treat population in phase III trials. In fact, in phase II or in phase III clinical trials, patients suffering from serious adverse events would probably receive co-medication and continue treatment. Therefore, this population is hypothetical but may exist. Conversely, one could be interested in the expectation of the marker given that the subjects had no DLT at a given time point, also called the partly conditional expectation.

Following Rouanet et al. (2017), if we assume that missingness for DLT is independent of the current marker value, then to obtain the expected activity under both the immortal cohort and given that subjects had no DLT we should estimate

$$E(Y_{ij}|X_{ij}, U_i, S_i > s - 1) = X_{ij}^T \beta + \gamma E(U_i|X_{ij}, S_i > s - 1) + E(r_{ij}|X_{ij}, U_i, S_i > s - 1).$$

In fact $E(r_{ij}|X_{ij}, U_i, S_i > s - 1)$ is negligible (Rouanet et al., 2017). For the immortal cohort $E(U_i|X_{ij}, S_i > s - 1) = 0$, whereas given that subjects had no DLT $E(U_i|X_{ij}, S_i > s - 1) \neq 0$. For the estimation of $E(U_i|X_{ij}, S_i > s - 1)$, McCulloch et al. (2016) proposed an approximation based on a log-linear time to event model, which does not lead to similar flexibility in the modeling of the variation of the hazard in time.

2.4 Simulation settings and evaluation criteria

2.4.1 Framework

In order to assess the performance of our design we conducted a set of extensive simulations, covering a wide variety of scenarios. Several aspects of our simulations matched the ovarian cancer trial, introduced in Section 2.2.1, such as time intervals,

CA 125 biomarker measurements, and censoring process. To that end, we simulated a trial of six doses (1.2, 1.4, 1.6, 1.8, 2.0, 2.2) and a total of 60 patients. Eight principal scenarios were investigated (Table 2.3). Scenarios covered cases where the OD and the MTD coincided, the OD was below the MTD, and finally all doses were highly toxic. The parameters used for data generation of these scenarios are provided in Table A.3. Some of these scenarios served for our sensitivity analysis (Tables A.4-A.5).

2.4.2 Time-to-DLT

For the time-to-DLT outcome, we considered a maximum of six treatment cycles. This number stemmed from the DLT-TARGETT group (Postel-Vinay et al., 2014), who showed that severe toxicities occurred quite often after cycle 1 and were almost null after cycle 6. We assumed that the hazard decreased over time (Postel-Vinay et al., 2011, 2014). Typically, the target toxicity level over one treatment cycle is between 20% and 33%, so over six treatment cycles setting the target at 40% was quite a realistic choice. In the principal scenarios, data was generated from the model of analysis. For the analysis model, the dose slope was fixed, across all scenarios, at $a_g = -3.29$ and for data generation we changed this specific value assuming a maximum change of 0.7 units. Parameter $a_g = -3.29$ provided satisfactory operating characteristics in capturing the monotonic dose-toxicity relationship. Administrative censoring was at cycle 6.

2.4.3 Activity

For the activity endpoint, following the Gynecological Cancer Intergroup guidelines (Rustin et al., 2004), we simulated CA 125 patient data at fixed time intervals with baseline values being at least twice the upper limit of the reference range, $[0, 35]$ units/ml. The marker was measured once a week, corresponding to three measurements per cycle, so participants could have a maximum of $n = 18$ measurements at the end of the trial.

Along with the time-to-DLT model, that gave a total of nine parameters to estimate in the joint model. One dose was considered “equally” active to the MTD if their minimum biomarker measurements did not differ by more than 20 units/ml, so $\zeta = 20$.

2.4.4 Missing data

Patients remained in the study until DLT occurrence, disease progression, consent withdrawal or completion of the 6 cycles of treatment. In fact, in phase I around 93% of study participants leave the study by the end of cycle 6 (Postel-Vinay et al., 2011). In our simulations all censoring takes place at the end of the corresponding cycle. The first cause of missingness was DLT occurrence. After a DLT patients did not contribute biomarker data. The second cause was disease progression and that led to information missing regarding the toxicity profile. Disease progression was defined with respect to Rustin et al. (2011) stating that: a) patients with elevated CA 125 pretreatment and normalization of CA 125 must show evidence of CA 125 greater than, or equal to, twice the upper limit of the reference range and b) patients with elevated CA 125 pretreatment, which never normalizes, must show evidence of CA 125 greater than, or equal to, twice the nadir value. These two missingness mechanisms can be seen as at random. We assumed that patients consent withdrawal, the third cause of missing data, was independent of both toxicity and activity and may occur with a 8% probability at each treatment cycle. Finally, we accounted for 7% intermittent missing responses for each biomarker measurement, that entailed missing data but not censoring.

2.4.5 Main analysis

A total of 2000 simulations were replicated for each scenario. Interest lay in evaluating the percentage of correct OD identification, which corresponded to the dose administered to the last patient in the trial and the mean number of patients allocated over the dose

range in the study. Furthermore, we explored the percentage of simulations where joint modeling could not be fitted. Convergence criteria for likelihood maximum were based on the gradient function, with an absolute error tolerance of 10^{-3} .

2.4.6 Sensitivity analysis

A sensitivity analysis was performed assessing the robustness of the proposed design, under various conditions including also different dose-activity and time-activity relationships than the ones investigated in the principal scenarios. First, we have simulated two scenarios (Table 2.4) where hazard was increasing at each successive cycle, inducing cumulative toxicities. For the cycle effect we assumed both a steep and a smooth slope, scenarios 4.1 and 4.2 respectively. Till now, we assumed that the time-activity relationship had a parabolic shape and that doses were ordered. However, we were interested in investigating the model behavior under various biomarker trajectories over time. To that end, we simulated one scenario in which the above relationship reached a plateau and another one where doses crossed in time, see Figures 2.5 a)-b) accordingly. Additionally, we explored scenarios (Table 2.5) with various degrees of misspecification; toxicity data was generated based on a logistic model, whereas for activity a saturated dose-biomarker relationship 2.6 was simulated, with a separate parameter for every dose level,

$$\begin{aligned} y_{ij(l)} &= \beta_0 + \beta_1 t_{ij}^2 + \beta_2 t_{ij} d_{(1)} + \beta_3 t_{ij} d_{(2)} + \beta_4 t_{ij} d_{(3)} + \beta_5 t_{ij} d_{(4)} \\ &+ \beta_6 t_{ij} d_{(5)} + \beta_7 t_{ij} d_{(6)} + U_i t_{ij} + r_{ij}. \end{aligned} \quad (2.6)$$

Analysis models were not modified.

The principal scenarios (Table 2.3) were also evaluated under the case of larger variances (Table A.4), smaller sample size and larger variances (Table A.5) and different distributions for the random effects between the two models (Table 2.5). We assumed that for the activity model, random effects followed a gamma distribution $U \sim \text{Gamma}(2, 2)$ and for the toxicity model, the standard normal distribution $U \sim N(0, 1)$. In order to

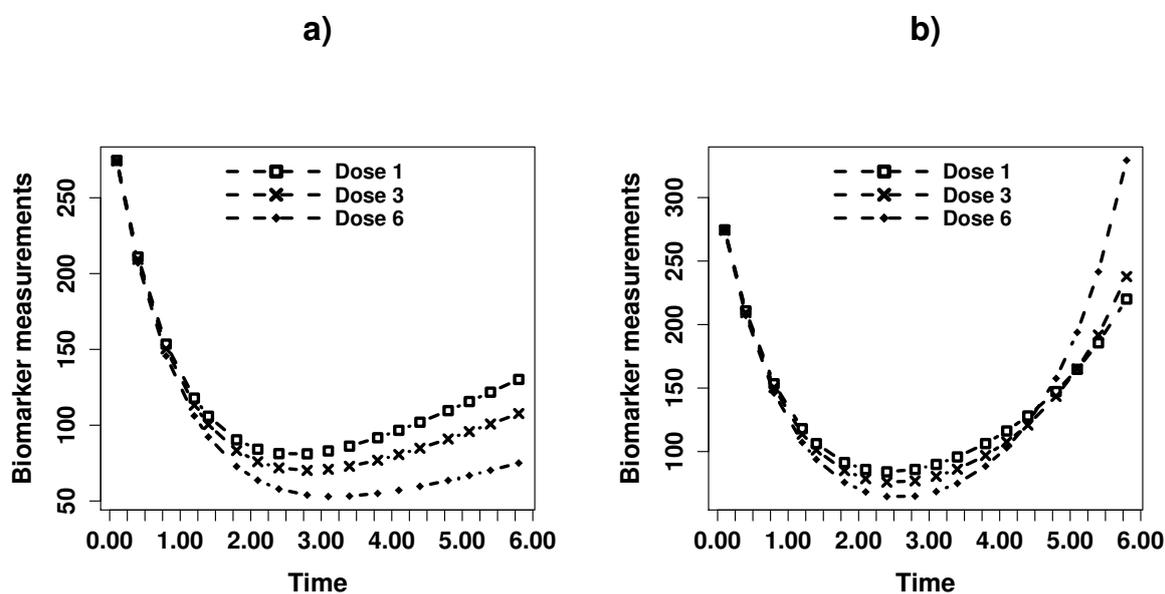


Figure 2.5: **Non-linear biomarker trajectories over time.**

assess which of the two models, toxicity or activity, has a stronger effect when being misspecified, we generated data where either one of the two models was modified. More specifically, for scenarios A.6.1-A.6.3 data was generated from a saturated linear mixed effects model for activity and the probit model for toxicity, whereas for scenarios A.6.4-A.6.6 data was generated from the linear mixed effects model for activity and a logistic model for toxicity (Table A.6). What is more, we considered a different biomarker-time relationship that corresponds to that depicted in Figure A.1. This trajectory refers to biomarkers, discussed in section 2.2.1, that reflect the direct action of the treatment such as the plasma concentration or the level of antibodies binding to their target. When patients respond to the treatment the levels of the biomarker increase with time and when they progress the levels decrease. We have explored our design under these types of biomarkers by simulating two additional scenarios (Table A.7). Data for the biomarker was generated from the same linear mixed effects model 2.2 that we introduced in sec-

tion 2.3.3. Scenarios A.7.1 and A.7.2 were generated in a such a way to match scenarios 5 and 2, respectively, from Table 2.3, at least in terms of toxicity and OD. Patients were censored due to lack of activity when their biomarker measurements decreased by more than 20 units/ml, after reaching the maximum activity. To further test the robustness of the joint model with one random effect, we replicated scenarios in which biomarker data was generated from a linear mixed effects model with both a random intercept, with variance σ_3^2 and a random slope and these two were highly correlated $\rho_{\sigma_1, \sigma_3}$ (Table A.8). Finally, we explored a case where activity data was generated assuming both a linear and a quadratic term for time (Table A.9). Scenario A.9.1 was created in such a way to match scenario 1 in Table 2.3, so as to allow for comparison of the two.

2.5 Results

2.5.1 Main results

Table 2.3 summarizes the main results. The percentage of correct selection for scenarios 1-7 was 97.3%, 79.5%, 98.4%, 90.9%, 99.5%, 94.8% and 86.5% accordingly. That was usually higher when the true OD was one of the extreme dose levels, so for instance scenarios 3, 5, and 6 and lower when the OD was in the middle of the dose range. The mean number of patients allocated to the OD ranged from 30 to 51 out of 60, depending on whether it coincided with the MTD. Regarding safety, a dose was recommended above the MTD for scenarios 1 and 3 in 1.5% and 0.2% of the simulations, accordingly. Generally, patients were not exposed to highly toxic doses, or to doses that could eventually be subtherapeutic. For scenarios 1-7, the median number of patients treated before the joint modeling initiation ranged from 17 to 28. For scenario 8, all doses were assumed to be highly toxic. The simulated trials came to an early halt in 83.3% of the simulations.

In our simulations, we observed a low rate of convergence failure, indicating that joint modeling could not be fitted. For scenarios 1,2,4,5, and 7, that rate was less than 0.44%, whereas for scenario 6 it slightly increased to 2.75%. Only in scenario 3, where dose 1 was

Table 2.3: Simulation results: Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). Activity data correspond to that presented in Figure 2.2. The optimal dose is in bold and the MTD in italic.

Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
1	$(Y_{(l),min}, p_l)$	(200, 0.00)	(176, 0.00)	(147, 0.02)	(115, 0.12)	(79, 0.38)	(40, 0.75)	
	$P_{\%}$	0.0	0.0	0.0	1.2	97.3	1.5	0.0
	\bar{N}_{pat}	2.0	2.0	2.2	5.4	40.5	7.9	
2	$(Y_{(l),min}, p_l)$	(71, 0.00)	(64, 0.00)	(57, 0.03)	(49, 0.12)	<i>(39, 0.38)</i>	(29, 0.75)	
	$P_{\%}$	0.2	11.0	79.5	8.9	0.4	0.0	0.0
	\bar{N}_{pat}	2.2	7.4	31.4	8.4	5.6	5.0	
3	$(Y_{(l),min}, p_l)$	(87, 0.32)	(78, 0.71)	(67, 0.96)	(55, 0.99)	(41, 0.99)	(26, 0.99)	
	$P_{\%}$	98.4	0.2	0.1	0.0	0.0	0.0	1.3
	\bar{N}_{pat}	51.3	6.5	2.0	0.2	0.0	0.0	
4	$(Y_{(l),min}, p_l)$	(87, 0.00)	(78, 0.00)	(67, 0.03)	(55, 0.12)	<i>(41, 0.38)</i>	(26, 0.75)	
	$P_{\%}$	0.0	0.0	3.8	90.9	5.3	0.0	0.0
	\bar{N}_{pat}	2.0	2.1	5.6	35.8	9.3	5.3	
5	$(Y_{(l),min}, p_l)$	(71, 0.02)	(68, 0.07)	(65, 0.20)	<i>(61, 0.43)</i>	(57, 0.68)	(52, 0.88)	
	$P_{\%}$	99.5	0.1	0.0	0.4	0.0	0.0	0.0
	\bar{N}_{pat}	30.0	3.6	6.3	11.4	7.4	1.3	
6	$(Y_{(l),min}, p_l)$	(200, 0.00)	(176, 0.00)	(147, 0.01)	(115, 0.04)	(79, 0.16)	(40, 0.41)	
	$P_{\%}$	0.0	0.0	0.0	0.0	5.2	94.8	0.0
	\bar{N}_{pat}	2.0	2.0	2.0	2.4	8.0	43.6	
7	$(Y_{(l),min}, p_l)$	(84, 0.00)	(73, 0.00)	(61, 0.01)	(48, 0.04)	(32, 0.16)	<i>(16, 0.41)</i>	
	$P_{\%}$	0.0	0.0	0.0	13.1	86.5	0.4	0.0
	\bar{N}_{pat}	2.0	2.0	2.2	9.0	34.2	10.6	
8	$(Y_{(l),min}, p_l)$	(71, 0.71)	(64, 0.88)	(57, 0.97)	(49, 0.99)	(39, 0.99)	(29, 0.99)	
	$P_{\%}$	16.7	0.0	0.0	0.0	0.0	0.0	83.3
	\bar{N}_{pat}	14.9	0.9	0.0	0.0	0.0	0.0	

both the OD and the MTD, joint model could not be fitted in 6.47% of the simulations. Thus, these simulations were conducted with the probit CRM exclusively. However, that could be expected as on average 57 patients were assigned to dose level 1, which made it harder to estimate parameters related to dose in the longitudinal model. Nonetheless, patients were still allocated to the correct OD.

2.5.2 Sensitivity analysis results

Results from the sensitivity analysis can be seen, first, in Table 2.4. For scenarios 4.1 and 4.2 we assumed a cumulative risk of toxicity. In case of cumulative toxicity, when no or low toxicity is observed in the first cycles and data is missing at later cycles due to censoring, estimating the true probability of toxicity can be challenging. We can see that by comparing the rate of OD identification of scenarios 4.1 and 2, of Table 2.3, that are almost identical. The power to detect a time-trend was 97% and 47% for scenarios 4.1 and 4.2, accordingly, illustrating the additional information that can be drawn from the sequential probit model regarding the risk of cumulative toxicity over time. The conditional risk of DLT on the MTD, for scenarios 4.1 and 4.2, increased from 0.01% to 21% and from 3% to 10%, respectively. An increasing hazard would probably be compatible with long drug administration. For scenarios 4.3 and 4.4, we simulated different biomarker trajectories over time (Figure 2.5 a)-b)). Our design is somehow sensitive when the time-activity relationship reaches a plateau, but not when doses cross over time (65% and 95.3% correct OD selection, respectively). In Table 2.5, there are five scenarios, for which data was generated from different models than the ones used for the analysis. The rate of correct OD identification for scenarios 5.1 to 5.5 was 74.8%, 54.2%, 99.2%, 97.4% and 92.7%, respectively. As expected, that rate was lower than the one from Table 2.3. It is worth mentioning, though, that for these simulations, doses were quite close in terms of probability of toxicity, which makes it harder for any algorithm to correctly identify the MTD and subsequently the OD.

The last part of our sensitivity analysis can be found in Appendix A. Increasing the

variance (Table A.4) or altering the random effects between the two models (Table A.5) did not affect the rate of correct OD selection. Increasing the variance and setting the sample size at 40 (Table A.5) was associated with a 7% lower rate of correct OD selection. Scenarios, for which either one of the two models was misspecified, are presented in Table A.6. Trials reach the true OD more frequently when the activity model is misspecified, rather than the toxicity, which is expected given the amount of information resulting from the repeated measurements. The percentage of correct OD selection was the same for scenarios A.7.1 and 5 from Table 2.3 and differed by 6% for A.7.2 and 2 (Table A.7). Thus, we conclude that there are no substantial differences stemming from the increasing or decreasing pattern of the biomarker over time. Assuming activity data with both random intercept and random slope did not affect the results, unless we increase the random intercept variance. More precisely, when the variance increased the correct OD selection drop from 80.3% to 67.9% (Scenarios A.8.1 and A.8.2 Table A.8). Finally, when activity data generation included both a linear and a quadratic term for time did not seem to impact the OD identification (Table A.9). In fact the observed rate in scenario A.9.1 is almost identical to that of scenario 1 in Table 2.3.

2.6 Discussion

In this chapter, we presented an adaptive design for the identification of the OD in phase I/II trials of MTAs in oncology. We considered a discrete time failure model to match the data usually collected in phase I and to evaluate modification of the hazard of toxicity over time as well as cumulative toxicities. Allowing for a longer DLT evaluation period, instead of per cycle modeling, could be an alternative but that would lead to patients dropping out. Nevertheless, in the case of patient reported outcomes, where the exact time of the DLT is known, the approach of Rizopoulos (2012) could be applied, assuming that approximation of the likelihood is feasible under small sample sizes. We believe in the importance of joint modeling for two major reasons. First, we tackled the issues of missing at random data, due to DLTs and lack of activity, as well as intermittent

Table 2.4: Sensitivity analysis of 2000 replicates and a sample size of 60. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). For scenarios 4.1 – 4.2 toxicity data was generated assuming increasing hazard at each successive treatment cycle. For scenarios 4.3 – 4.4 data was generated for different biomarker trajectories over time, as shown in Figure 2.5 a) and Figure 2.5 b), respectively. The optimal dose is in bold and the MTD in italic.

Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
4.1	$(Y_{(t),min}, p_t)$	(71, 0.00)	(64, 0.01)	(57, 0.03)	(49, 0.14)	<i>(39, 0.38)</i>	(29, 0.72)	
	$P_{\%}$	10.7	14.6	64.2	9.7	0.8	0.0	0.0
	\bar{N}_{pat}	5.2	8.0	19.8	6.6	3.7	16.7	
4.2	$(Y_{(t),min}, p_t)$	(168, 0.00)	(151, 0.00)	(132, 0.00)	(112, 0.03)	(89, 0.14)	(63, 0.42)	
	$P_{\%}$	0.8	0.0	0.1	0.0	6.4	92.7	0.0
	\bar{N}_{pat}	2.2	2.0	2.1	2.4	8.1	43.2	
4.3	$(Y_{(t),min}, p_t)$	(81, 0.04)	(76, 0.13)	<i>(70, 0.35)</i>	(65, 0.65)	(59, 0.89)	(53, 0.98)	
	$P_{\%}$	65.0	17.7	16.8	0.5	0.0	0.0	0.0
	\bar{N}_{pat}	19.7	12.7	17.5	8.3	1.7	0.1	
4.4	$(Y_{(t),min}, p_t)$	(84, 0.00)	(80, 0.00)	(76, 0.03)	(72, 0.12)	<i>(68, 0.38)</i>	(65, 0.75)	
	$P_{\%}$	95.3	0.9	0.0	0.0	3.8	0.0	0.0
	\bar{N}_{pat}	34.8	5.2	3.3	4.1	7.3	5.3	

Table 2.5: Sensitivity analysis of 2000 replicates and a sample size of 60. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). Data generation from a saturated linear mixed effects model for activity and a logistic model for toxicity. The optimal is in bold and the MTD in italic.

Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
5.1	$(Y_{(t),min}, p_t)$	(86, 0.02)	(73, 0.04)	(69, 0.08)	(66, 0.13)	(64, 0.22)	<i>(58, 0.34)</i>	
	$P_{\%}$	20.8	74.8	3.8	0.6	0.0	0.0	0.0
	\bar{N}_{pat}	8.9	27.1	7.0	3.9	3.8	9.3	
5.2	$(Y_{(t),min}, p_t)$	(119, 0.09)	(113, 0.14)	(88, 0.23)	<i>(80, 0.35)</i>	(65, 0.51)	(46, 0.69)	
	$P_{\%}$	8.0	18.8	54.2	12.3	6.2	0.5	0.0
	\bar{N}_{pat}	5.5	10.4	22.4	10.8	7.6	3.3	
5.3	$(Y_{(t),min}, p_t)$	(79, 0.22)	<i>(77, 0.33)</i>	(74, 0.48)	(72, 0.65)	(70, 0.80)	(66, 0.91)	
	$P_{\%}$	99.2	0.5	0.2	0.0	0.0	0.0	0.1
	\bar{N}_{pat}	44.4	6.6	5.1	2.8	0.9	0.2	
5.4	$(Y_{(t),min}, p_t)$	(79, 0.02)	(77, 0.04)	(74, 0.07)	(72, 0.12)	(70, 0.19)	<i>(66, 0.30)</i>	
	$P_{\%}$	97.4	0.0	0.0	0.0	0.2	2.4	0.0
	\bar{N}_{pat}	34.5	3.1	3.3	3.6	4.0	11.5	
5.5	$(Y_{(t),min}, p_t)$	(175, 0.02)	(160, 0.04)	(138, 0.07)	(109, 0.12)	(84, 0.19)	<i>(53, 0.30)</i>	
	$P_{\%}$	1.0	0.0	0.0	0.2	6.1	92.7	0.0
	\bar{N}_{pat}	2.4	2.2	2.4	3.1	7.8	42.1	

missing responses. Second, patients in the study were allocated to the OD that is not necessarily the MTD, as it would be the case if only toxicity data was considered. Our principal goal was to integrate all the information collected during the trial, without being restricted to the first treatment cycle only and without loss of information due to dichotomization of the activity measurements.

The design was evaluated on diverse settings, showing a percentage of correct OD selection particularly high and good operating characteristics. Most importantly, simulated subjects were rarely exposed to doses above the MTD and a large fraction of them was not undertreated. Even under convergence failure of the joint modeling, with the aid of CRM, patients were still administered to safe and highly active doses.

A limitation of our model is associated to the type of plateau it can capture. When either of the dose-activity or time-activity relationships reaches a steep plateau, our design might be less robust. Strictly speaking, this refers to the case where biomarker measurements on the plateau have exactly the same value and are not just “equally” active as per our definition. Under this scenario, for the dose-activity relationship, for instance, the OD would be located on the beginning of the plateau, since with further escalation only toxicity would be gained. Even though determination of the MTD does not pose a problem, the model cannot properly identify the OD, and as a corollary the recommended dose at the end of the trial is usually closer to the MTD. In practice, that does not constitute an issue, since the dose selected is between the true OD and the MTD, meaning that it has the maximum possible activity, and yet is safe.

This specific design is quite flexible and can be applied to any biomarker that has a monotonic or plateau type relationship with the dose, irrespectively of the increasing or decreasing pattern over time. It can be adapted by researchers to include as many cycles and biomarker measurements as needed, and additionally to allow for non-balanced data. For the decision process, alternative criteria can be considered, such as a toxicity-activity trade-off (Thall and Cook, 2004). Overall, we consider that the inclusion of multiple and continuous biomarker measurements may play a pivotal role in integrating all the information gathered during a clinical trial.

Chapter 3

Optimal dose selection considering both toxicity and activity data; plateau detection for molecularly targeted agents

3.1 Introduction

In the previous chapter, we presented an adaptive design for the identification of the OD. We explored several scenarios, showing that our method is robust, under diverse settings. At the end, we discussed a limitation of the design regarding the case of a steep plateau, in which doses would have exactly the same biomarker value and consequently exactly the same activity. In that case, our method was not very efficacious in identifying at which dose is located the beginning of the plateau. Although the algorithm could still propose a dose as “active” as the MTD, we may need a method that can identify the beginning of the plateau. A major benefit of such an approach would be for cases where many doses have exactly the same activity and the MTD is far from the OD. Identification of the beginning of the plateau could help select a dose “equally” active to MTD and yet a lot less toxic.

Furthermore, a significant part of our adaptive design relies on the joint modeling

technique. In the previous chapter, we mentioned that several authors have shown in simulation studies the importance of implementing joint models, when two outcomes are believed to be associated. This is due to the fact that joint models reduce the bias, leading to more robust estimations. However, an important question is whether reducing the estimation bias through joint modeling could help improve the dose selection, as compared to implementing independent modeling, i.e. considering independently the activity and the toxicity model. Cunanan and Koopmeiners (2014) performed a simulation study, investigating whether operating characteristics of a phase I/II clinical trial improved after the implementation of joint modeling. For their study, they used a copula model, where both toxicity and activity were considered as binary endpoints. They found that both approaches performed just as well due to difficulty in estimating the copula model correlation parameters from binary data. Nonetheless, there is merit in examining if this is the case in the presence of continuous and repeated biomarker measurements that are richer in information and under the likelihood inference.

In this chapter, we propose an extension of the adaptive design described in Chapter 2. We rely on the same joint modeling technique of longitudinal continuous biomarker measurements and time-to-DLT data, using a shared random effect and likelihood approach. We use the same model for toxicity and a change-point model for activity. Specifically, for the biomarker measurements, we employ a linear mixed effects model that can account for either a plateau or a linear dose-activity relationship (Figure 3.1). Every time we fit the joint model, the algorithm selects the activity model (plateau or linear) that best fits the observed data. Similar to the previous design, the MTD is the dose with a cumulative risk of DLT, over six cycles that minimizes the distance from a predefined TTL. The OD is defined as the lowest dose, within a range of highly active doses, that is below or equal to the beginning of the plateau, if the beginning of the plateau is below the MTD. If the MTD is a dose below the plateau beginning, then the OD is the lowest dose, within a range of highly active doses, that is below or equal to the MTD. The motivating example of our design is the multicenter phase IB trial in patients with platinum resistant epithelial ovarian carcinoma, described in Section 2.2.1. Finally, we investigated the contribution of joint modeling versus independent modeling in the case

of a phase I/II clinical trial.

The remainder of Chapter 3 is organized as follows. In Section 3.2, we introduce the models for toxicity and activity, the study design and the decision process for the identification of the OD. In Section 3.3, we describe the simulation framework and the scenarios that were investigated. In Section 3.4, we comment on the results of the simulation study and we compare this design with the one described in Chapter 2. Section 3.4.3 is devoted to the results obtained from the comparison of joint modeling with independent modeling. Discussion follows in Section 3.5.

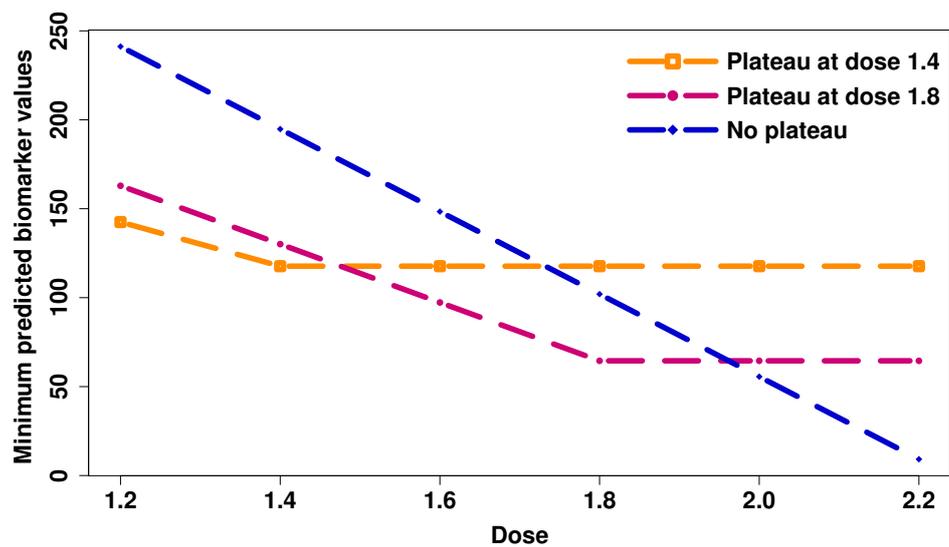


Figure 3.1: Dose-biomarker relations.

3.2 Methods

3.2.1 General framework

First, we shall briefly recap some basic notation of Chapter 2. Consider a trial of N individuals, $i = 1, \dots, N$ and a total of M planned doses to be tested $D = \{d_{(1)}, \dots, d_{(M)}\}$. Patients are sequentially assigned to a specific dose level, selected based on the data from previous patients. Cycles of treatment are repeatedly administered until DLT, disease progression, defined by some elevation of the biomarker, or consent withdrawal. At the end of each cycle study participants are evaluated as to whether they had a DLT during the cycle or not. Therefore, for the time-to-DLT we consider a discrete time scale, $C = \{1, \dots, k\}$, where k is the total number of treatment cycles. As before, we consider a total of six treatment cycles $k = 6$. The biomarker levels are measured at each weekly visit, in order to determine if and when patients respond to the treatment and when they progress. For the biomarker measurements, we have continuous measurements on a continuous time scale $\mathcal{T} = (0, k)$. If a DLT is observed or the agent is not active, based on the biomarker data, then the patient is removed from the study, resulting in censored DLT data and MAR for the biomarker.

3.2.2 Modeling time-to-DLT

Let S_i be the time-to-DLT. In this design, for the time-to-DLT outcome we consider the same probit model 2.3,

$$P(S_i = s | S_i > s - 1, U_i) = 1 - \Phi(a_0 + a_1 c_{i(s-1)} + a_g d_i + \gamma U_i), \quad c_{i(s-1)} \in \mathcal{C},$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution, a_0 and a_1 are unknown parameters, d_i is the patient's dose and $c_{i(s-1)}$ is the administered cycle. Once again we fix the slope of the dose to a constant value a_g . $U_i \sim N(0, \sigma_1^2)$ is the shared random effect

that comes from the activity model and γ is the corresponding parameter for estimation. We remind that the probability of having a DLT at cycle s is conditional on having no DLT in all the past cycles.

The probability of DLT for a given patient at each cycle is obtained as follows:

$$P_{(l)}(S_i = s|U_i) = \begin{cases} (1 - \Phi_{(l)1}) & s = 1 \\ \Phi_{(l)1}\Phi_{(l)2} \dots \Phi_{(l),s-1}(1 - \Phi_{(l)s}) & s = 2, \dots, 6 \\ \Phi_{(l)1}\Phi_{(l)2} \dots \Phi_{(l),s-1}\Phi_{(l)s} & s > 6, \end{cases}$$

where $1 - \Phi_{(l)s}$ is the hazard function at cycle s and dose level l , l being the index for the dose. Then, the cumulative probability of DLT at the end of cycle 6 is

$$P_{(l)}(S_i \leq 6|U_i) = \sum_{s=1}^6 P_{(l)}(S_i = s|U_i).$$

3.2.3 Modeling activity

Let $Y_{ij(l)}$ be the repeated continuous biomarker measurements for the i^{th} individual, at times $t_{ij} \in \mathcal{T}$, $j = 1, \dots, n$ represents the visit number, and at dose level l . Now, we shall consider an activity model different from the one presented in Chapter 2. The model structure, regarding the time-biomarker relationship remains the same, motivated by the activity trend over time, shown in Figure 2.2. However, for the dose-biomarker relationship we allow for more flexibility. Let pl be an integer that takes values from 1 to $M - 1$, indicating the dose level that is the beginning of the plateau and 1 be the indicator function. Then, doses of treatment modify the course of CA 125 as follows

$$y_{ij(l)} = \beta_0 + \beta_1 t_{ij}^2 + \beta_2 t_{ij} (d_{i(l)} 1(l < pl) + d_{i(pl)} 1(l \geq pl)) + U_i t_{ij} + r_{ij}, \quad (3.1)$$

where β_0 , β_1 , and β_2 are unknown parameters and $d_{i(l)}$ is the dose. U_i refers to the random effect associated with the time of the longitudinal measurements and R_{ij} are

the model residuals that are mutually independent and follow a Gaussian distribution, $MVN(\mathbf{0}, \sigma_2^2 I)$, with I being the identity matrix. Model 3.1 makes the assumption that dose modifies the biomarker trajectory linearly and after a certain dose level the biomarker reaches a plateau. This plateau is observed at a specific time point for each patient and this time can vary among patients. Therefore, when the dose level is lower than pl , the dose-biomarker relationship monotonically decreases, and when the dose level is equal to or higher than pl , the dose-biomarker relationship plateaus. This model is similar to the one proposed by Riviere et al. (2016). Doses on the plateau have exactly the same biomarker value. In case where the dose-activity relationship is strictly monotonic and there is no plateau model 3.1 becomes

$$y_{ij(l)} = \beta_0 + \beta_1 t_{ij}^2 + \beta_2 t_{ij} d_{i(l)} + U_i t_{ij} + r_{ij}. \quad (3.2)$$

An example is illustrated in Figure 3.1. The orange and the violet curves are two examples of model 3.1, in which the dose-activity relationship monotonically decreases and then plateau begins at dose level 2 and 4, respectively. On the other hand, the blue curve delineates the monotonic relationship, described by model 3.2. Hence, for the 6 dose levels pl takes values between 1 and 5. Value 1 is for the monotonic relationship, whereas values 2 to 5 are the possible dose levels that could indicate the beginning of a plateau. Every time we apply the joint modeling, we investigate all of the 5 candidate models. In particular, we estimate each of the five joint models for the different dose-activity relationships and we select the one that best fits the data using the Akaike information criterion (AIC) (Akaike, 1973). Let ξ_θ be the number of parameters we estimate for each model. Then taking the likelihood 2.4, described in Section 2.3.4, the AIC is

$$AIC = 2 * \xi_\theta - (2 * (-\log(L(\boldsymbol{\theta}; y, s, \delta))).$$

We select the joint model with the minimum value for the AIC.

3.2.4 Study design

The study design is very similar to the one described in Section 2.3.5. As before we split the design into two stages. For the first one dose escalation proceeds based on toxicity only and for the second one, we identify the OD based on both outcomes. We start with a ‘2+2’ design, where cohorts of patients escalate, de-escalate or remain at the same dose based on the DLTs observed in the first treatment cycle. When two DLTs and two non-DLTs have occurred, we stop the ‘2+2’. If all 3 criteria, 1) at least 16 patients have been enrolled, 2) the standard errors of all joint model parameter estimates are below 20, and 3) the Hessian matrix is positive definite, are satisfied, we switch to the joint modeling stage. If not, then we implement the probit time-to-DLT CRM 2.5 to guide dose-escalation. We continue with the CRM until all 3 criteria are satisfied and then we switch to the joint modeling. Recall Figure 3.2 that depicts the study design procedure.

Finally, when all 3 criteria are met, we initiate the second stage and we update activity and toxicity estimates after each new patient, using joint modeling. It is in the second stage, where we consider both toxicity and activity data that this design defers from the previous one. This is due to the fact that every time we apply the joint modeling, we investigate several models for the biomarker measurements and we are not restricted to a single one. Based on the model selected with AIC, we determine the MTD and the OD that is the dose the new patient will be allocated to. We repeat this procedure every time there is a new patient in the study. The model selected after each new patient is not necessarily the same. Nevertheless, after a certain number of patients has been enrolled in the trial, we expect it to eventually converge to one of the five candidate models. As before, for safety reasons dose skipping is not allowed.

Finally, we implemented the same early stopping rule for excessive toxicity at dose level 1, proposed by Ivanova et al. (2005). We generated a Pocock-type boundary for repeated testing of the toxicity at the first treatment cycle. A probability of DLT above 35% at the 10% level (one-sided) led to stop the trial.

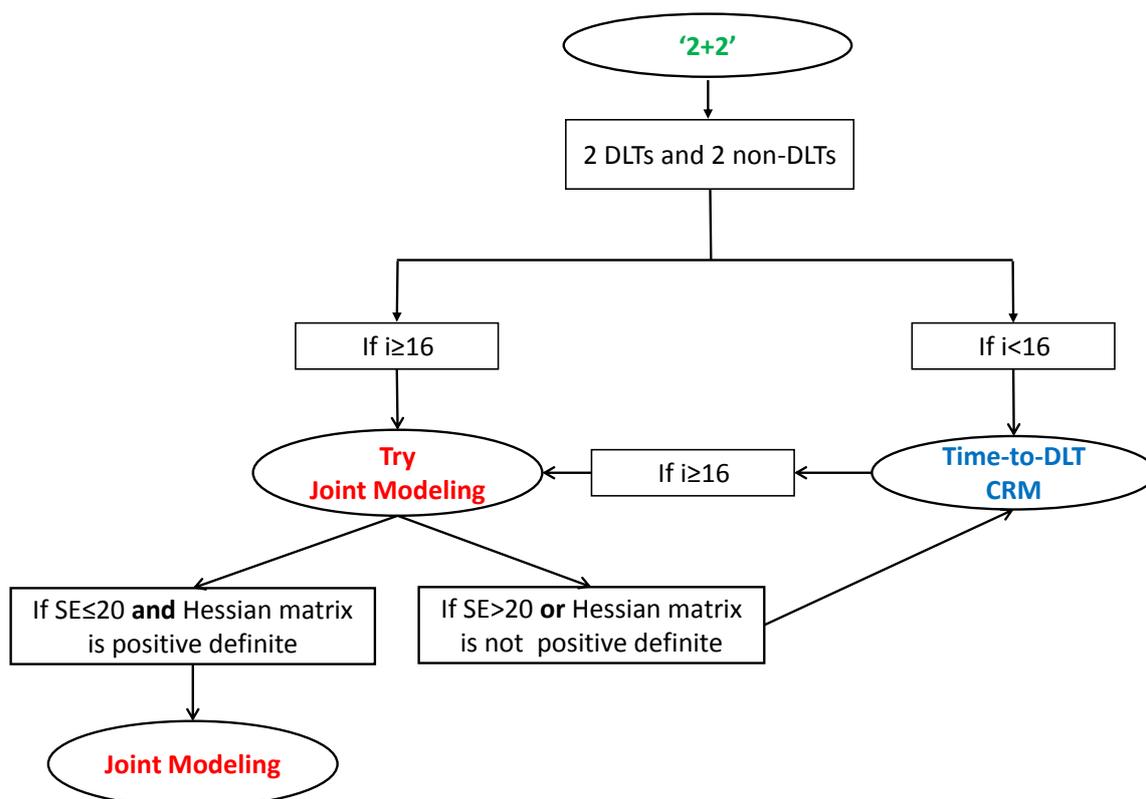


Figure 3.2: Representation of the study design described in Section 2.3.5. Abbreviations: DLTs, Dose Limiting Toxicities; i , index for the patient; SE, Standard Error; CRM, Continual Reassessment Method.

3.2.5 Decision process

Based on the joint model estimates, we determine the MTD and the OD. Similarly to our previous adaptive design, we start from the MTD, which is defined cumulatively over six cycles. Taking the conditional probability for $U = 0$ we minimize over all dose levels

$$MTD = \min_{(l)} |\epsilon - P_{(l)}(S \leq 6 | U = 0)|,$$

where ϵ is the TTL. Then the MTD is used to select a subset of safe doses B . This subset contains the MTD and all doses below the MTD, if the MTD is not the first dose level.

The OD is defined as the lowest dose, within a range of highly active doses. We remind that if the biomarker measurements reach a plateau, then all doses located on that plateau are considered to have exactly the same biomarker activity. Therefore, for the selection of the OD we identify two cases; 1) the MTD is a dose higher than the dose indicating the beginning of the plateau and 2) the MTD is a dose equal or lower to the beginning of the plateau.

For the first case, the OD is defined as the lowest dose, within a range of highly active doses, below or equal to the dose indicating the beginning of the plateau. Define \mathcal{M} to be a subset of doses below or equal to the plateau beginning and $\mathcal{M} \subseteq B$. Then, for every $d_{(l)} \in \mathcal{M}$ we estimate, at each time visit t_j and for each dose level l , the conditional predicted biomarker activity

$$\hat{y}_{j(l)} = \hat{\beta}_0 + \hat{\beta}_1 t_j^2 + \hat{\beta}_2 t_j (d_{(l)} 1(l < pl) + d_{(pl)} 1(l \geq pl)). \quad (3.3)$$

As before, the maximum drug activity is associated with the minimum biomarker measurement over time, thus, for each dose we select the minimum of the predicted values $\hat{y}_{min,(l)}$. We identify doses that are “equally” active to the dose located on the beginning of the plateau as follows,

$$|\hat{y}_{min(pl)} - \hat{y}_{min(l)}| \leq \zeta,$$

where ζ is the equivalence range. The lowest dose in \mathcal{M} is selected for the next patient. An example of such a case is illustrated in Figure 3.3 a). In this example, the MTD is dose level 6 and the beginning of the plateau is located at dose level 4. Imagine, for the purpose of this example and without loss of generality that the minimum biomarker measurement of every dose is at the same time visit, as shown in Figure 3.3 a). Then, $\Omega = \{d_{(1)}, d_{(2)}, d_{(3)}, d_{(4)}\}$ and $\zeta = 20$ that is any dose $d_{(l)} \in \mathcal{M}$ to be “equally” active to the plateau, if $|\hat{y}_{min(pl)} - \hat{y}_{min(l)}| \leq 20$. For instance, for dose level 3, we have

$|\hat{y}_{min(pl)} - \hat{y}_{min(3)}| = 17$, meaning that $d_{(3)}$ is “equally” active to $d_{(4)}$ that corresponds to the beginning of the plateau. Therefore, the OD is $d_{(3)}$.

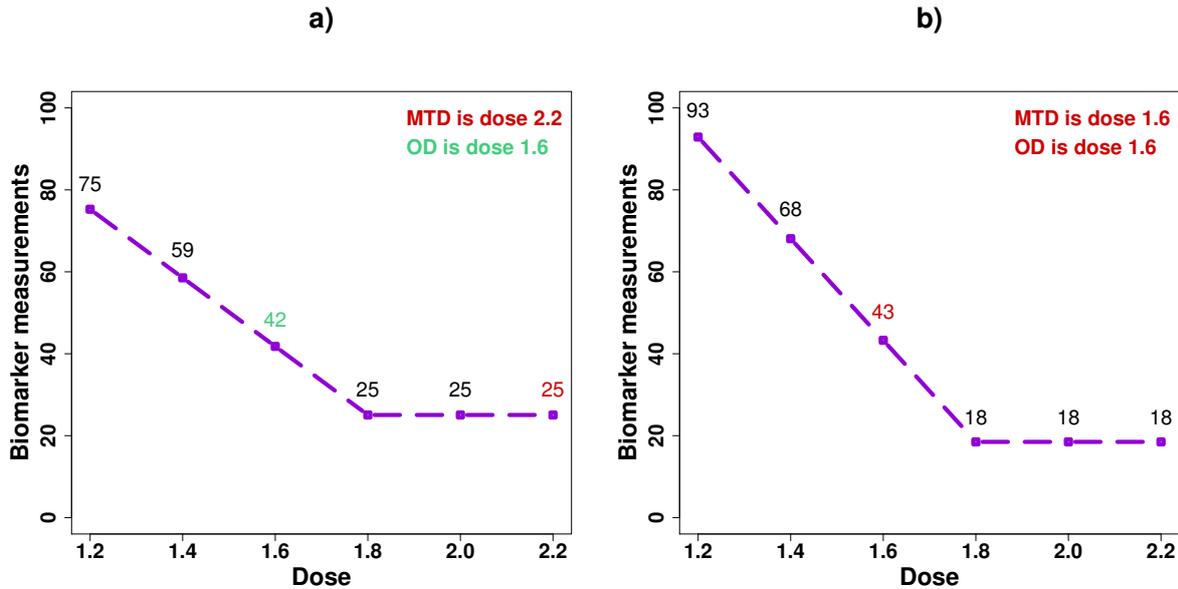


Figure 3.3: Dose-activity relationships.

For the second case, the OD is defined as the lowest dose, within a range of highly active doses, below or equal to the MTD. Then, for every $d_{(l)} \in B$ we estimate, at each time visit t_j and for each dose level l , the predicted biomarker activity, using model 3.3. We identify doses that are “equally” active to the MTD as follows,

$$|\hat{y}_{min(MTD)} - \hat{y}_{min(l)}| \leq \zeta.$$

Define Λ to be a subset including all doses that are “equally” active to the MTD. The lowest dose in Λ is selected for the next patient. Figure 3.3 b) is an example of such a

scenario. Here, the MTD is dose level 3 and the beginning of the plateau is located at dose level 4. Then, $B = \{d_{(1)}, d_{(2)}, d_{(3)}\}$ and as earlier $\zeta = 20$. In this example the MTD is also the OD, since for dose level 2 we have $|\hat{y}_{\min(MTD)} - \hat{y}_{\min(2)}| = 25$, meaning that $d_{(2)}$ is not “equally” active to the MTD.

Estimation of the MTD and the OD is repeated after every new patient. When the MTD is a dose higher than the start of the plateau, this OD definition is slightly different from that described in Section 2.3.6. When we consider the case of plateau, the two different decision process algorithms could potentially lead to different results. An example of such a case is discussed in Section 3.4.2.

3.2.6 Independent modeling

For the independent modeling approach, we considered the same toxicity and activity models that were presented in sections 3.2.2 and 3.2.3. For the toxicity outcome and after removing the random effects we get

$$P(S_i = s | S_i > s - 1) = 1 - \Phi(a_0 + a_1 c_{i(s-1)} + a_g d_i), \quad c_{i(s-1)} \in \mathcal{C}.$$

For the activity outcome model 3.1 was used.

3.3 Simulation settings and evaluation criteria

3.3.1 Framework

To assess the performance of this design, we conducted a set of extensive simulations, covering a wide variety of scenarios. The simulation framework was similar to that presented in Section 2.4, so as to compare this design with the one described in Chapter 2.

Scenarios covered cases where the OD, the MTD, and the beginning of the plateau coincided, the OD was below the MTD and the beginning of the plateau, and finally all doses were highly toxic.

We rapidly remind the main characteristics of the simulated trials. We assumed a trial of six doses (1.2, 1.4, 1.6, 1.8, 2.0, 2.2) and a total of 60 patients. For the toxicity endpoint, we considered a maximum of six treatment cycles. We supposed that the risk of DLT decreased at each treatment cycle and that the TTL after 6 cycles was 40%. For the analysis model, the dose slope was fixed, across all scenarios, at $a_g = -3.29$. For the activity endpoint, following the Gynecological Cancer Intergroup guidelines (Rustin et al., 2004), we simulated CA 125 patient data at fixed time intervals with baseline values being at least twice the upper limit of the reference range, $[0, 35]$ units/ml. The marker was measured once a week, corresponding to three measurements per cycle, so participants could have a maximum of $n = 18$ measurements at the end of the trial. As before, one dose was considered “equally” active to the MTD or to the beginning of the plateau if their minimum biomarker measurements did not differ by more than 20 units/ml, so $\zeta = 20$.

In our simulations all censoring takes place at the end of the corresponding cycle. The first cause of missingness was DLT occurrence. After a DLT patients did not contribute data anymore on the biomarker. The second cause was disease progression and that led to information missing regarding the toxicity profile. Censoring due to disease progression was applied with respect to Rustin et al. (2011). Administrative censoring was at cycle 6. We assumed that patients consent withdrawal, the third cause of missing data, was independent of both toxicity and activity and may occur with a 8% probability at each treatment cycle. Finally, we accounted for 7% intermittent missing responses for each biomarker measurement, that entailed missing data but not censoring.

3.3.2 Main analysis

For the main analysis, we investigated five principal scenarios (Table 3.1). A total of 1000 simulations were replicated for each scenario. Interest lay in evaluating the percentage of correct OD identification, which corresponded to the dose administered to the last patient in the trial and the mean number of patients allocated over the dose range in the study. Another aspect under investigation was the percentage of simulations where joint modeling could not be fitted. Particularly, we assessed whether the use of a more flexible activity model could help resolve convergence issues observed in the previous design. Convergence criteria for likelihood maximum were based on the gradient function, with an absolute error tolerance of 10^{-3} . The parameters used for data generation of these scenarios are provided in Table B.1. For the principal scenarios, data was generated from the model of analysis.

3.3.3 Sensitivity analysis

In order to assess the robustness of the proposed design, a sensitivity analysis was performed, accounting for various degrees of model misspecification. To begin with, we explored scenarios in which toxicity data was generated based on a logistic model (Table 3.2). For the activity, data was simulated from the saturated dose-biomarker relationship 2.6, presented in Section 2.4.6, with a separate parameter for every dose level. Additionally, we compared this design with the previous one, by simulating activity data with the log model 2.2 (Table 3.3). More specifically, for data generation we selected four from the eight principal scenarios, investigated in Table 2.3, and we analyzed them with the plateau model. The aim was first to investigate the model behavior under model misspecification and second to examine if the percentage of convergence failure would diminish, given the model flexibility.

Next, the principal scenarios (Table 3.1) were also evaluated under the case of larger variances with smaller sample size and smaller sample size with different random effects distributions between the two models (Table B.2). In particular, we assumed that for

the activity model, random effects followed a gamma distribution $U \sim \text{Gamma}(2, 2)$ and for the toxicity model, the standard normal distribution $U \sim N(0, 1)$. Furthermore, we investigated two scenarios (Table B.3), where hazard was increasing at each successive cycle, inducing cumulative toxicities. For the cycle effect we assumed both a steep and a smooth slope, scenarios B.3.1 and B.3.2, respectively. In the preceding simulations, we had accounted for scenarios, in which biomarker data was generated from a linear mixed effects model with both a random intercept and a random slope. Here, we considered this setting and additionally, activity data was generated from the log model 2.2. The investigated scenarios (Table B.4) were exactly the same as those presented in Table A.8, allowing for direct comparison of the two approaches.

Analysis models were not modified.

3.3.4 Comparison of joint modeling and independent modeling

In order to compare the joint modeling with the independent modeling approach, we generated two different types of scenarios. For the first one, we assumed that simulated data was correlated with an increasing degree of correlation between the toxicity and the activity model (Figure 3.5 and Figure 3.6). Given the correlated data, we varied the sample size, the residual and random effect variance and additionally, generated toxicity data from the previously described logistic model. For the second type of scenarios, we assumed that toxicity and activity data was not correlated, $\gamma = 0$, and similarly, we investigated the model behavior under diverse sample sizes, variances and misspecified models (Figure 3.7 and Figure 3.8). Data was analyzed with both approaches, so as to determine whether joint modeling is more robust than independent modeling concerning correct dose selection. Since for the OD selection the first step is to correctly identify the MTD, the main focus of these simulations was the percentage of correct MTD selection.

3.4 Results

3.4.1 Main results

Table 3.1 summarizes the main results. The percentage of correct OD selection for scenarios 1-4 was 99.2%, 99.2%, 99.7%, and 86.3% accordingly. The mean number of patients allocated to the OD ranged from 27 to 43 out of 60, depending on whether it coincided with the MTD. As far as safety is concerned, no dose was ever recommended at the end of the trial that was above the MTD. Therefore, patients were not exposed to highly toxic doses. We can also see that most patients were allocated to doses close to the MTD or the OD, an indication that they were not exposed to subtherapeutic doses either. For scenarios 1-4, the median number of patients treated before the joint modeling initiation ranged from 18 to 28. These numbers were almost identical to those presented in Section 2.5.1, for scenarios 1-7 of Table 2.3. For scenario 5, all doses were assumed to be highly toxic. The simulated trials came to an early halt in 81.9% of the simulations.

In the Chapter 2, we discussed the issue of convergence failure, i.e. the cases where the joint model could not be fitted. For scenarios 1-4 the failure rate was 0% in all our simulations. Therefore, all simulations proceeded with the joint modeling. We believe that this could be the result of the more flexible activity model and the fact that at each new patient entry the algorithm selects the activity model that best fits the data, without being restricted to a single model.

3.4.2 Sensitivity analysis results

In Table 3.2, there are four scenarios, for which data was generated from different models than the ones used for the analysis. The rates of correct OD identification for scenarios 2.1 to 2.4 were 73.4%, 78.5%, 99.5%, and 53.3%, respectively. As expected, these rates were lower than the ones from Table 2.3. Nonetheless, that was expected, since doses were quite close to each other, in terms of probability of toxicity and simulated models

Table 3.1: Main analysis: Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). The optimal dose is in bold, the MTD in italic and the beginning of the plateau is underlined.

Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
1	$(Y_{(l),min}, p_l)$	(92, 0.02)	(68, 0.07)	(40, 0.20)	<i>(40, 0.43)</i>	(40, 0.68)	(40, 0.88)	
	$P_{\%}$	0.0	0.8	99.2	0.0	0.0	0.0	0.0
	\bar{N}_{pat}	2.1	2.8	34.7	11.3	7.7	1.4	
2	$(Y_{(l),min}, p_l)$	(140, 0.00)	(118, 0.00)	(92, 0.03)	(64, 0.12)	<i>(32, 0.38)</i>	(32, 0.75)	
	$P_{\%}$	0.1	0.0	0.0	0.7	99.2	0.0	0.0
	\bar{N}_{pat}	2.0	2.0	2.2	5.8	42.7	5.2	
3	$(Y_{(l),min}, p_l)$	(78, 0.07)	(74, 0.19)	<i>(69, 0.39)</i>	(63, 0.64)	(57, 0.84)	(50, 0.96)	
	$P_{\%}$	99.7	0.3	0.0	0.0	0.0	0.0	0.0
	\bar{N}_{pat}	31.8	6.0	11.0	9.1	2.0	0.1	
4	$(Y_{(l),min}, p_l)$	(96, 0.02)	(88, 0.07)	(80, 0.20)	<i>(71, 0.43)</i>	(71, 0.68)	(71, 0.88)	
	$P_{\%}$	13.6	86.3	0.1	0.0	0.0	0.0	0.0
	\bar{N}_{pat}	6.4	26.9	6.3	11.2	7.8	1.4	
5	$(Y_{(l),min}, p_l)$	(78, 0.71)	(74, 0.88)	(69, 0.97)	(63, 0.99)	(57, 0.99)	(50, 0.99)	
	$P_{\%}$	18.1	0.0	0.0	0.0	0.0	0.0	81.9
	\bar{N}_{pat}	17.5	1.1	0.0	0.0	0.0	0.0	

were misspecified. It was previously pointed out that due to small differences in the dose finding algorithm, it is possible that the two approaches (Chapter 2 and Chapter 3) can identify different doses as the OD. An example of such a case is scenario 2.1. With the previous design, the OD would be dose level 2, since we would be interested in a difference of 20 units/ml from the MTD, which is dose level 6. However, with the new design, there is a plateau that starts at dose level 2 and therefore, the OD was dose level 1 that did

not differ by more than 20 units/ml from the plateau. Due to the model misspecification, it is not clear where the plateau begins and for that reason the dose-activity relationship of scenario 2.1 is illustrated in Figure 3.4.

Table 3.2: Sensitivity analysis of 1000 replicates and a sample size of 60. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). Data generation from a saturated linear mixed effects model 2.6 for activity and a logistic model for toxicity. The optimal dose is in bold, the MTD in italic and the beginning of the plateau is underlined.

Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
2.1	$(Y_{(l),min}, p_l)$	(86, 0.03)	<u>(73, 0.04)</u>	(69, 0.08)	(66, 0.13)	(64, 0.22)	<i>(58, 0.34)</i>	
	$P_{\%}$	73.4	19.8	6.8	0.0	0.0	0.0	0.0
	\bar{N}_{pat}	23.7	11.7	8.0	3.5	3.8	9.3	
2.2	$(Y_{(l),min}, p_l)$	(119, 0.09)	(113, 0.14)	(88, 0.23)	<i>(80, 0.35)</i>	(65, 0.51)	(46, 0.69)	
	$P_{\%}$	8.9	2.9	78.5	1.4	8.3	0.0	0.0
	\bar{N}_{pat}	5.5	10.4	22.4	10.8	7.6	3.3	
2.3	$(Y_{(l),min}, p_l)$	(79, 0.22)	<i>(77, 0.33)</i>	(74, 0.48)	(72, 0.65)	(70, 0.80)	(66, 0.91)	
	$P_{\%}$	99.5	0.0	0.0	0.0	0.0	0.0	0.5
	\bar{N}_{pat}	45.6	5.5	4.6	3.0	1.1	0.2	
2.4	$(Y_{(l),min}, p_l)$	(175, 0.02)	(160, 0.04)	(138, 0.07)	(109, 0.12)	(92, 0.19)	<i>(74, 0.30)</i>	
	$P_{\%}$	1.6	0.2	3.2	53.3	41.7	0.0	0.0
	\bar{N}_{pat}	2.8	2.5	3.9	26.6	14.0	10.2	

Table 3.3 presents the results of the comparison of the two designs. Even though the activity data was simulated from the log model, the plateau model gave the same percentage of correct OD selection in all four scenarios. Most importantly, for scenario

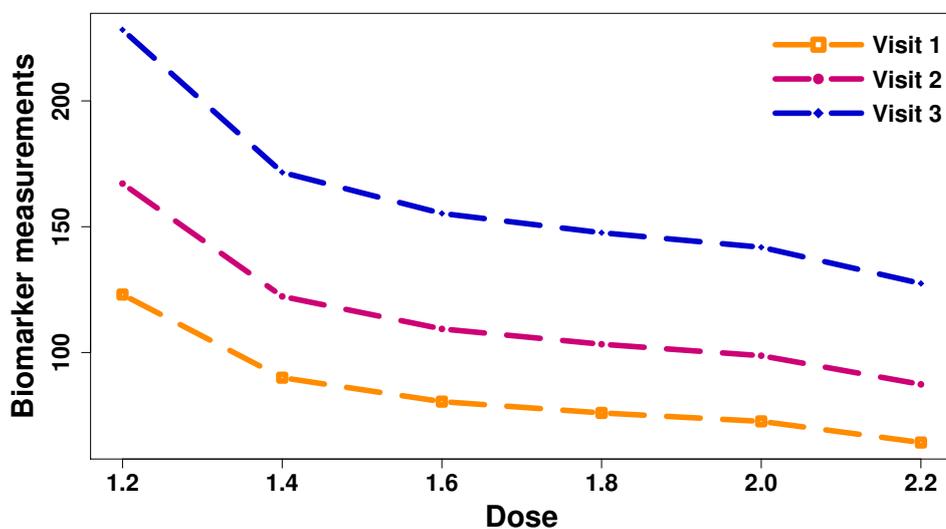


Figure 3.4: Dose-biomarker relationship for three different time visits. This relationship was generated by model 2.6.

3.2, the log model had 6.47% convergence failure, i.e. these simulations were conducted with the probit CRM alone. On the contrary, for the plateau model, all simulations proceeded with the joint modeling, meaning that there was 0% convergence failure.

The remaining of the sensitivity analysis can be found in Appendix B. Increasing the variance and setting the sample size at 40 (Table B.2) was associated with a small change in the percentage of OD identification, compared to Table 3.1. More precisely, correct OD selection was between 2% – 7% lower for these scenarios. Similarly, altering the random effects distribution for the activity and toxicity model and setting again the sample size at 40 (Table B.2) did not seem to have a significant impact on the results. Correct OD selection was 0% – 5% lower for these scenarios than those of Table 3.1. For scenarios B.3.1 and B.3.2 (Table B.3), we assumed a cumulative risk of toxicity. The percentage of correct OD selection for the two scenarios was 100% and 93.5%. For the toxicity outcome, we selected similar scenarios to those presented in Table 2.4. Surprisingly, this design had a higher rate of OD identification when the OD was in the middle of the dose range (scenario B.3.1 versus 4.1). When the OD was on the tail of the dose range

both designs produced similar results. Finally, in Table B.4, we have two more scenarios comparing the two designs. If we assume activity data with both a random intercept and a random slope, and simulated from the log model, both designs perform well. However, if we increase the variance the plateau model seems to outperform the log model (67.9% versus 73.9%). The good performance of the plateau model, under the case of larger variances, could be related to the model flexibility.

3.4.3 Comparison of joint modeling and independent modeling

Figures 3.5 and 3.6 summarize the simulation results, assuming that the activity and toxicity data is correlated. If we assume small variance and 60 patients (Figure 3.5 a)), the two approaches perform relatively the same way and when the correlation increases, $\gamma \geq 0.6$, the joint modeling approach has a 6% higher rate of correct MTD identification. When increasing both the residual and the random effect variance (Figure 3.5 b)) joint modeling has a better performance in correctly selecting the MTD that ranges from 11% to 18%, when the correlation increases from 0.4 to 0.8. Further increase of the random effect variance (Figure 3.5 c)), leads to a better performance of the joint modeling approach ranging from 11% up to 21%, when the correlation increases from 0.2 to 0.8. Finally, after increasing the overall variance and setting the sample size at 40 patients, joint modeling outperforms the independent modeling approach by 7% to 11%, when the correlation increases from 0.6 to 0.8.

Figure 3.6 illustrates similar scenarios to those presented in Figure 3.5, with the difference that data for the toxicity model come from a logistic model. Overall, the two approaches perform similarly. Joint modeling outperforms independent modeling by 8% to 11% (Figure 3.6 b), d)) when increasing the residual and random effect variance.

Figures 3.7 and 3.8 summarize the simulation results, assuming that the activity and toxicity data is not correlated. Under the case of increasing variance and irrespectively of the sample size, 60 patients (Figure 3.7 a)) or 40 patients (Figure 3.7 b)), the two approaches produce identical results. Finally, the only case we identified, where indepen-

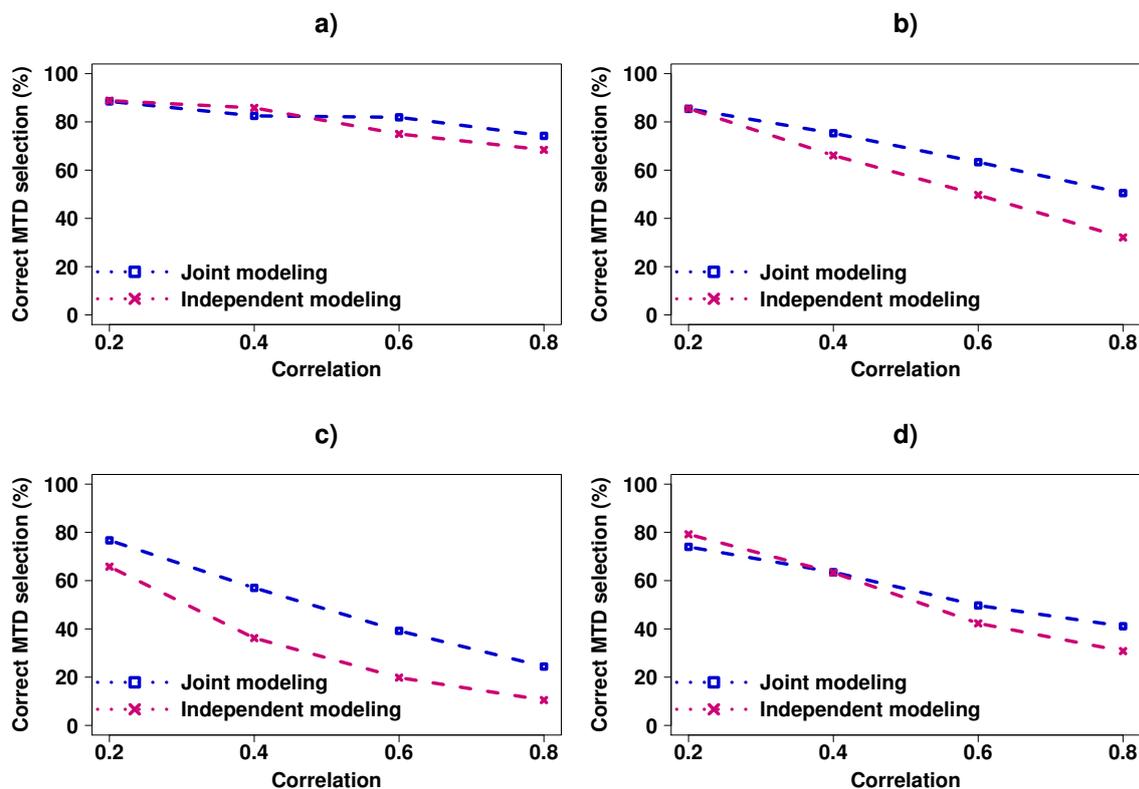


Figure 3.5: Percentage of correct MTD selection in the presence of increasing correlation between the activity and toxicity data. a) Variance ($\sigma_1 = 1$ and $\sigma_2 = 3$) and $N=60$ patients. b) Variance ($\sigma_1 = 2$ and $\sigma_2 = 5$) and $N=60$ patients. c) Variance ($\sigma_1 = 4$ and $\sigma_2 = 3$) and $N=60$ patients. d) Variance ($\sigma_1 = 2$ and $\sigma_2 = 5$) and $N=40$ patients.

dent modeling outperforms the joint modeling by 6%, is under model misspecification and small overall variance ((Figure 3.8 a))).

3.5 Discussion

In this chapter, we proposed an extension of the adaptive design for the identification of the OD, presented in Chapter 2. The principal goal was to improve the previous

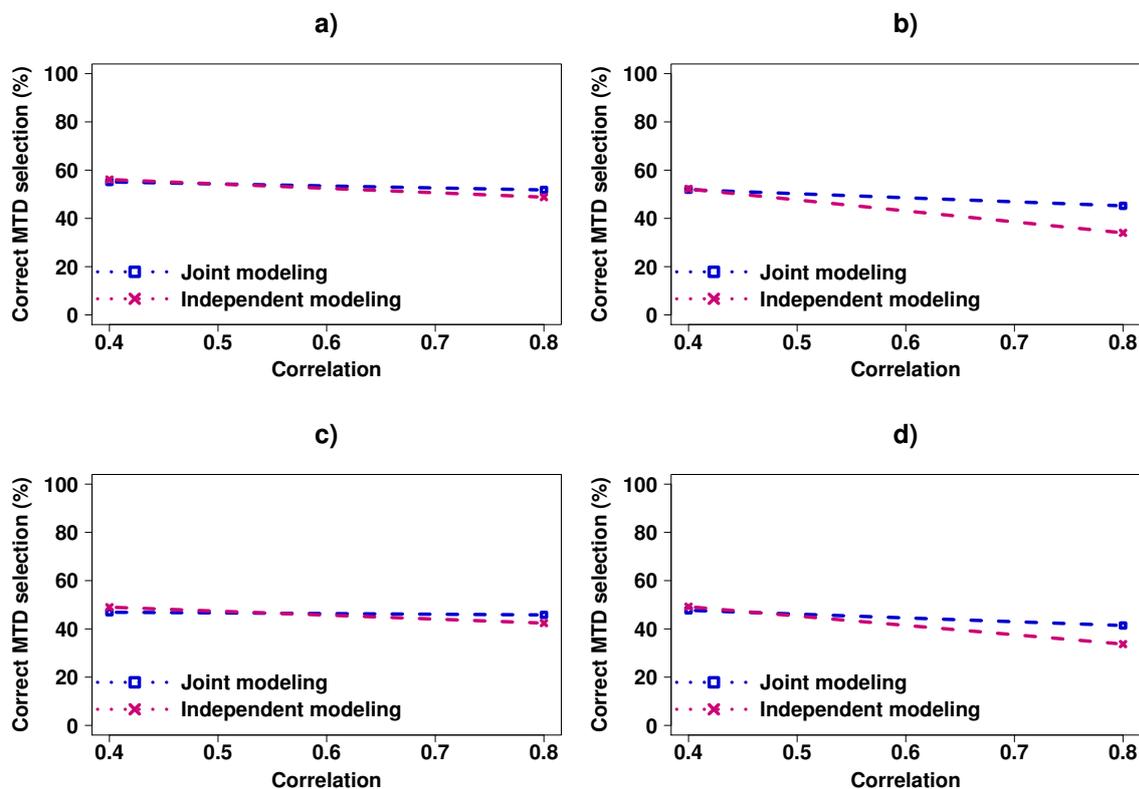


Figure 3.6: Percentage of correct MTD selection in the presence of increasing correlation between the activity and toxicity data. Toxicity data was generated under a logistic model. a) Variance ($\sigma_1 = 1$ and $\sigma_2 = 3$) and $N=60$ patients. b) Variance ($\sigma_1 = 2$ and $\sigma_2 = 5$) and $N=60$ patients. c) Variance ($\sigma_1 = 2$ and $\sigma_2 = 5$) and $N=40$ patients. d) Variance ($\sigma_1 = 2$ and $\sigma_2 = 5$) and $N=40$ patients.

design, mainly in terms of dose-activity relationship. As before, for the toxicity outcome, we considered a discrete time failure model in order to evaluate the modification of the hazard, during the course of a phase I trial, as well as the cumulative toxicities. For the biomarker measurements, we employed a linear mixed effects model that allowed for both monotonic and plateau dose-activity relationships, but not parabolic ones. Furthermore, we addressed the important case of missing at random data, due to DLTs and lack of activity, as well as intermittent missing responses. Finally, we investigated the impact of using joint modeling in the presence or not of correlated toxicity and activity data, versus independent modeling.

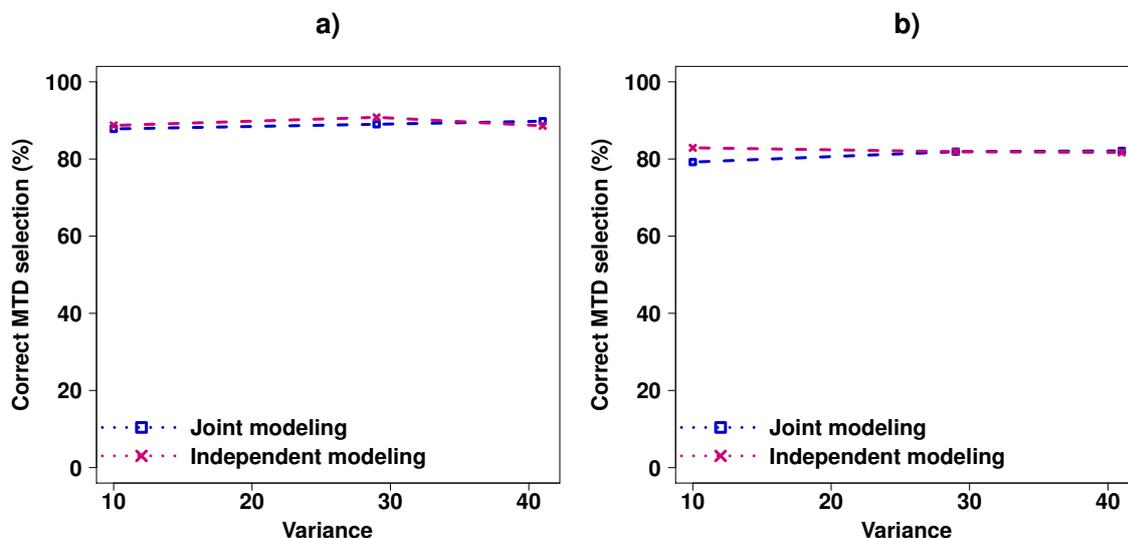


Figure 3.7: Percentage of correct MTD selection in the presence of increasing variance. Absence of correlation between the activity and toxicity data. a) N=60 patients. b) N=40 patients.

This design was assessed under various settings and was also compared to the precedent design. We observed a high rate of correct OD selection and study participants were most often not exposed to highly toxic or to subtherapeutic doses. The structure of the activity model, due to its flexibility, allowed to capture several biomarker trajectories, including the one of the log model. Most importantly, because of selecting the model that best fits the data, we were able to tackle the issue of convergence failure and to conduct all simulations with the joint modeling algorithm. Comparing the two approaches showed that in most of the scenarios both designs performed equally good and in one scenario the plateau model outperformed the log. Nonetheless, the two designs were not always directly comparable, due to the nature of the dose finding algorithm.

In theory, any linear model could be applied for the activity outcome. Increasing the number of the investigated models could help increase the robustness of the design, since it would be feasible to select a model closer to the observed data. Likewise, for the dose finding algorithm it could be possible to restrict the OD selection on the dose located on the beginning of the plateau and remove the definition of “equally” active doses. Finally,

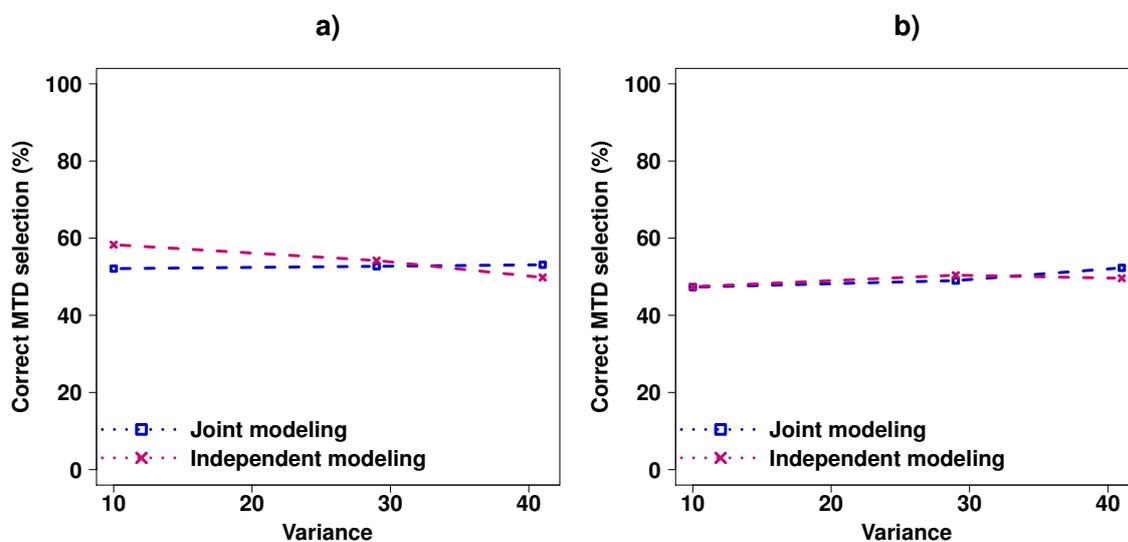


Figure 3.8: Percentage of correct MTD selection in the presence of increasing variance. Absence of correlation between the activity and toxicity data. Toxicity data was generated under a logistic model. a) N=60 patients. b) N=40 patients.

alternative dose escalation schemes, for the trial initiation, could be considered.

The performance of joint modeling versus independent modeling was evaluated under certain settings, showing that in the presence of large variances, the former outperforms the latter. Even in the absence of correlated toxicity and activity data the two approaches performed similarly. Hence, we conclude that the estimation of the additional association parameter γ does not impede the correct MTD identification, but on the contrary, in the presence of correlation it can enhance it.

Table 3.3: Sensitivity analysis of 1000 replicates and a sample size of 60. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). Data was generated from the models described in Chapter 2. The optimal dose is in bold and the MTD in italic.

Model		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
3.1	$(Y_{(l),min}, p_l)$	(71, 0.00)	(64, 0.00)	(57, 0.03)	(49, 0.12)	<i>(39, 0.38)</i>	(29, 0.75)	
log	$P_{\%}$	0.2	11.0	79.5	8.9	0.4	0.0	0.0
	\bar{N}_{pat}	2.2	7.4	31.4	8.4	5.6	5.0	
plateau	$P_{\%}$	2.0	12.1	79.5	3.4	3.0	0.0	0.0
	\bar{N}_{pat}	3.0	7.5	31.6	5.7	7.2	7.2	
3.2	$(Y_{(l),min}, p_l)$	(87, 0.32)	(78, 0.71)	(67, 0.96)	(55, 0.99)	(41, 0.99)	(26, 0.99)	
log	$P_{\%}$	98.4	0.2	0.1	0.0	0.0	0.0	1.3
	\bar{N}_{pat}	51.3	6.5	2.0	0.2	0.0	0.0	
plateau	$P_{\%}$	98.7	0.1	0.0	0.0	0.0	0.0	1.2
	\bar{N}_{pat}	51.4	5.9	1.9	0.1	0.0	0.0	
3.3	$(Y_{(l),min}, p_l)$	(71, 0.02)	(68, 0.07)	(65, 0.20)	<i>(61, 0.43)</i>	(57, 0.68)	(52, 0.88)	
log	$P_{\%}$	99.5	0.1	0.0	0.4	0.0	0.0	0.0
	\bar{N}_{pat}	30.0	3.6	6.3	11.4	7.4	1.3	
plateau	$P_{\%}$	100.0	0.0	0.0	0.0	0.0	0.0	0.0
	\bar{N}_{pat}	29.9	3.5	6.1	11.0	7.9	1.5	
3.4	$(Y_{(l),min}, p_l)$	(84, 0.00)	(73, 0.00)	(61, 0.01)	(48, 0.04)	(32, 0.16)	<i>(16, 0.41)</i>	
log	$P_{\%}$	0.0	0.0	0.0	13.1	86.5	0.4	0.0
	\bar{N}_{pat}	2.0	2.0	2.2	9.0	34.2	10.6	
plateau	$P_{\%}$	0.2	0.0	0.0	13.6	86.2	0.0	0.0
	\bar{N}_{pat}	2.1	2.0	2.2	9.8	34.1	9.8	

Chapter 4

Impact of cumulative toxicity in phase I trials. What should we expect at the recommended phase II dose?

4.1 Introduction

In Chapters 1, 2, and 3 we discussed the importance of considering time, in dose-finding methods, for the selection of the MTD and the OD and we presented designs to that end. What we called OD in the previous chapters, here it will be referred to as RP2D. As mentioned earlier the EMA (European Medicines Agency, 2016) pointed out the significance to improve phase I designs of MTAs, by changing the definition of the MTD, allowing for a broader DLT-evaluation period and finally considering AEs of several treatment cycles for the RP2D assessment. This requirement is even stronger with immune-toxic side effects, for which the median time varies from 5 to 15 weeks, which is beyond the usual DLT assessment period (3-4 weeks) (Champiat et al., 2016). In Chapters 2 and 3 we used a definition of the MTD cumulatively over 6 treatment cycles assuming a TTL of 40%. However, whilst for guiding dose-escalation, the notion of what constitutes an acceptable rate of acute toxicity is rather well-defined, when it

comes to guiding the RP2D recommendation, no such definition of what constitutes an acceptable cumulative and per-cycle rate of toxicity, in single-agent administration, exists. Considering cumulative toxicity is even of greater importance when the per-cycle risk of toxicity increases, indicating that the repeated administration of an agent could potentially be life-threatening. In particular, when 20% of acute toxicity is observed at the MTD, what cumulative toxicity rate should we expect over several treatment cycles?

Another major aspect when evaluating the MTD and the RP2D is related to the type of toxicity we observe. Two large classes of toxicity are the hematologic and the non-hematologic ones. These toxicities do not have the same impact on the human body. More precisely, hematologic toxicities have a shorter duration, are less hazardous and are usually more easily treated than the non-hematologic ones. Therefore, it is important to consider them separately since the acceptable risk of AEs per-cycle and cumulatively for these two classes of toxicity is not necessarily the same.

The remainder of Chapter 4 is organized as follows. In Section 4.2, we present the objective of this work. Section 4.3 is the methods part, where we introduce our data, the models and the various analyses. In Section 4.4 we describe our data and the results of our analyses, complemented by figures and tables. Discussion follows in Section 4.5.

4.2 Objective

The aim of the present work is to provide an overview of the risk of first-severe toxicity per treatment cycle and of the corresponding cumulative incidence function (CIF) over up to six treatment cycles. We estimated these risks based on 26 phase I clinical trials of MTAs administered as single agents, from the Cancer Therapy Evaluation Program (CTEP) of the US NCI. A secondary objective is to document the relation between time-on-treatment and the risk of severe toxicity. We provide a nomogram that relates the risk of severe toxicity at cycle 1 with estimations of the CIF expected over up to six cycles of treatment. Finally, we investigate the CIFs of non-hematologic and hematologic severe toxicity, within the competing risks framework.

4.3 Methods

4.3.1 Trials and patients characteristics

This retrospective analysis included single agent phase I studies that reached the MTD and were provided to the DLT-TARGETT group. The DLT-TARGETT group is a European Organisation for Research and Treatment of Cancer (EORTC)-led initiative. The aim was to provide a comprehensive description of all drug-related toxicities reported in phase I trials of MTAs (in terms of type, grade, cycle of occurrence) and to ultimately propose recommendations for DLT definition and phase II dose recommendation process, customized for phase I trials of novel MTAs. Collected trials dated from 1997 to 2013. Information recorded included treatment arms, agent under investigation, method of treatment administration, duration of treatment cycle, the MTD, etc. All adult patients with solid tumors or lymphomas who received at least one cycle of treatment were eligible for the analysis. Individual data, provided to the DLT-TARGETT group, recorded for each patient the number of treatment cycles, the dose administered, the follow up, observed toxicities per cycle and grade, and DLTs, as defined per study protocol, among others. Administered doses were measured in different units among studies, therefore, we standardized them, by dividing the amount of the dose by the corresponding trial's MTD.

4.3.2 Toxicity data

For the analysis we considered data up to six treatment cycles. As for Chapters 2 and 3, this number was motivated by Postel-Vinay et al. (2014), who showed that severe toxicities occurred quite often after cycle 1 and were almost null after cycle 6. All AEs at least potentially related to the treatment that were not present at baseline at the same or higher grade were extracted. In the beginning of the thesis we call attention to the fact that DLTs are defined for the first treatment cycle only. DLTs are usually associated to severe toxicities of grade 3, 4 or 5. Thus, since DLTs are not documented after cycle 1, the outcome of interest is severe toxicity of grade 3,4, or 5. Toxicity severity was harmonized across studies using the NCI Common Terminology Criteria of Adverse Events, version 3.0. Severe toxicities were further divided into hematologic and non-hematologic, according to Medical Dictionary for Regulatory Activities 15.

4.3.3 Main statistical analysis

A treatment cycle as defined per protocol was used as time unit, irrespectively of the duration in days. As for our previous analyses, we considered a discrete time scale with a maximum of 6 cycles $C = \{1, \dots, 6\}$. The event of interest was the time-to-first severe toxicity (grade 3, 4 or 5). The per-cycle risk of having an AE was estimated for those still at risk at the cycle initiation (i.e. the hazard function). For all analyses we relied on a probit model similar to that introduced in Section 2.3.2. Let S_i be the time-to-first severe toxicity, for the i^{th} individual, at time s given that no toxicity occurred at time $s - 1$. Then,

$$P(S_i = s | S_i > s - 1) = 1 - \Phi(a_0 + a_1 c_{i(s-1)} + a_2 d_i), \quad c_{i(s-1)} \in \mathcal{C}, \quad (4.1)$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution, a_0 , a_1 , and a_2 are unknown parameters, d_i is the dose administered to the i^{th} patient and $c_{i(s-1)}$ is the cycle variable. For the estimation of the CIF of severe toxicity over up to 6 cycles, as in Section 2.3.2, we summed up the risk of toxicity at each cycle

$$P(S_i \leq 6) = \sum_{s=1}^6 P(S_i = s).$$

For both cycle and dose we considered a linear relationship on the probit scale. To investigate if the linearity assumption was correct, we assessed the model residuals. Specifically, for the treatment cycle we investigated two different models, for patients treated at the MTD. In the first one, we assumed a different hazard for cycles 1-3 and 4-6 and in the second one a different hazard for cycles 1-2 and 3-6. The underline model is

$$P(S_i = s | S_i > s - 1) = 1 - \Phi(a_0 + a_1 c_{1i(s-1)} + a_2 c_{2i(s-1)}),$$

where c_1 and c_2 are the corresponding cycle variables. Similarly for the dose, we assumed a different dose effect for patients treated below the MTD and those treated at or above the MTD,

$$P(S_i = s | S_i > s - 1) = 1 - \Phi(a_0 + a_1 c_{i(s-1)} + a_2 d_i 1_{(l < MTD)} + a_3 d_i 1_{(l \geq MTD)}),$$

where l is the dose level and 1 is the indicator function.

In a first analysis, three subgroups of patients treated at doses below, above and at the MTD were created. The per-cycle risk of severe toxicity and the CIFs were estimated separately in all three groups. Then a complete model 4.1 adjusted on the dose was built to develop a nomogram predicting the CIF over six cycles, if the risk of severe toxicity observed at cycle 1 ranged from 5% to 35%. The same model was also estimated separately on the individual studies that were large enough to fit the model (17 / 26 studies).

In a second analysis, interest lay in evaluating the CIFs of hematologic and non-hematologic severe AEs. Hematologic toxicities included also mixed cases, i.e. hematologic and non-hematologic AEs observed on the same treatment cycle. Thus, there were strictly non-hematologic AEs and hematologic (with or without concomitant non-hematologic toxicity). First, we performed a more “naive” analysis, supposing that only one of the two risks was of interest and the other one was not considered. Therefore, we estimated the risk of having a hematologic toxicity, possibly in the presence of non-hematologic toxicities and vice versa. For a more elaborate analysis, we turned to the competing risks framework (Lee et al., 2018). In that case, we were interested in the risk of having a first hematologic or a first non-hematologic toxicity. All analyses were performed for the three dose subgroups, i.e. above, below, and at the MTD. For the estimation of the cause-specific hazards we applied model, similar to model 4.1. For instance for competing event (1), i.e. the non-hematologic severe toxicity, we have

$$P_{(1)}(S_i = s | S_i > s - 1) = 1 - \Phi(a_0 + a_1 c_{i(s-1)}), \quad c_{i(s-1)} \in \mathcal{C}.$$

Then, for the non-hematologic severe toxicity at cycle 6, the CIF is

$$\begin{aligned}
 P_{(1)}(S_i \leq 6) &= P_{(1)}(S_i = s | S_i > s - 1) \sum_{t=1}^{s-1} P(S_i > t - 1) \\
 &= P_{(1)}(S = 1) + P_{(1)}(S = 2 | S > 1)P(S > 1) + P_{(1)}(S = 3 | S > 2)P(S > 2) \\
 &+ P_{(1)}(S = 4 | S > 3)P(S > 3) + P_{(1)}(S = 5 | S > 4)P(S > 4) \\
 &+ P_{(1)}(S = 6 | S > 5)P(S > 5).
 \end{aligned}$$

When considering more than one treatment cycles we expect to have missing data. This data is the result of severe toxicities, disease progression or consent withdrawal. In our data, we had the information for severe toxicities, but we had no knowledge regarding the cause of censored observations. A property of the underline model 4.1 is that it takes into account censoring when estimating the risk of severe toxicity. Nonetheless, it makes the assumption that censored observations would have the same risk as the observations still at risk, which may not necessarily be true. Administrative censoring was considered at the end of cycle 6.

Prediction intervals were obtained from the bias corrected bootstrap technique (Efron, 1981).

4.3.4 Sensitivity analysis

As a first sensitivity analyses, we repeated the competing risks analysis classifying AEs into strictly hematologic and non-hematologic (with or without hematologic). For the second sensitivity analysis we excluded grade 3 hematologic toxicities from the definition of severe toxicities. As mentioned before, hematologic toxicities are often easier to treat by the clinicians, thus quite often grade 3 is not considered in the definition of the DLT. Then, we repeated the analyses for the three dose subgroups as well as the competing risks analysis. Competing risks analysis was conducted with the same definition of competing events as for the main analysis.

4.4 Results

4.4.1 Trial characteristics

From the 27 phase I studies provided by the NCI, one did not reach the MTD and was excluded from analysis (Figure 4.1). The 26 eligible trials enrolled a total of 942 patients with solid tumors or lymphomas. AEs data was available for all patients. There were 7.7% antiangiogenic agents, 3.8% antivascular agents, 7.7% CDK inhibitors, 11.5% HDAC inhibitors, 23.1% HSP inhibitors, 11.5% immunotherapy, 11.5% monoclonal antibodies, 7.7% proteasome inhibitors and 15.4% other classes of agents. Administration route was mainly intravenous (80.8%), followed by oral (15.4%) and intraperitoneal (3.9%). Cycle duration ranged from 14 to 42 days (Table 4.1).

4.4.2 Treatment administration

Of the 942 patients, 58.2% received a second cycle, 25.1%, 16.5%, 9.9%, and 7.2% received a 3rd to 6th cycle, respectively (Table 4.2). A total of 289 of the 942 patients (30.7%) were assigned to the dose later defined as the trial MTD. Of those 289, 20 (6.9%) received 6 cycles (Figure 4.2 a), Table C.1). A total of 490 patients (52%) and 163 (17.3%) were treated below and above the MTD respectively.

4.4.3 Toxicity outcomes

Over the 6 treatment cycles, 333 of the 942 patients (35.3%) had at least one severe toxicity of any type. Among these, 203 had only severe non-hematologic toxicity, 91 had solely severe hematologic toxicity and 39 had both severe non-hematologic and hematologic toxicity at the same cycle (Figure 4.2 b)). Non-hematologic toxicities included 39% gastrointestinal disorders, 10% general disorders, 7% central nervous system disor-

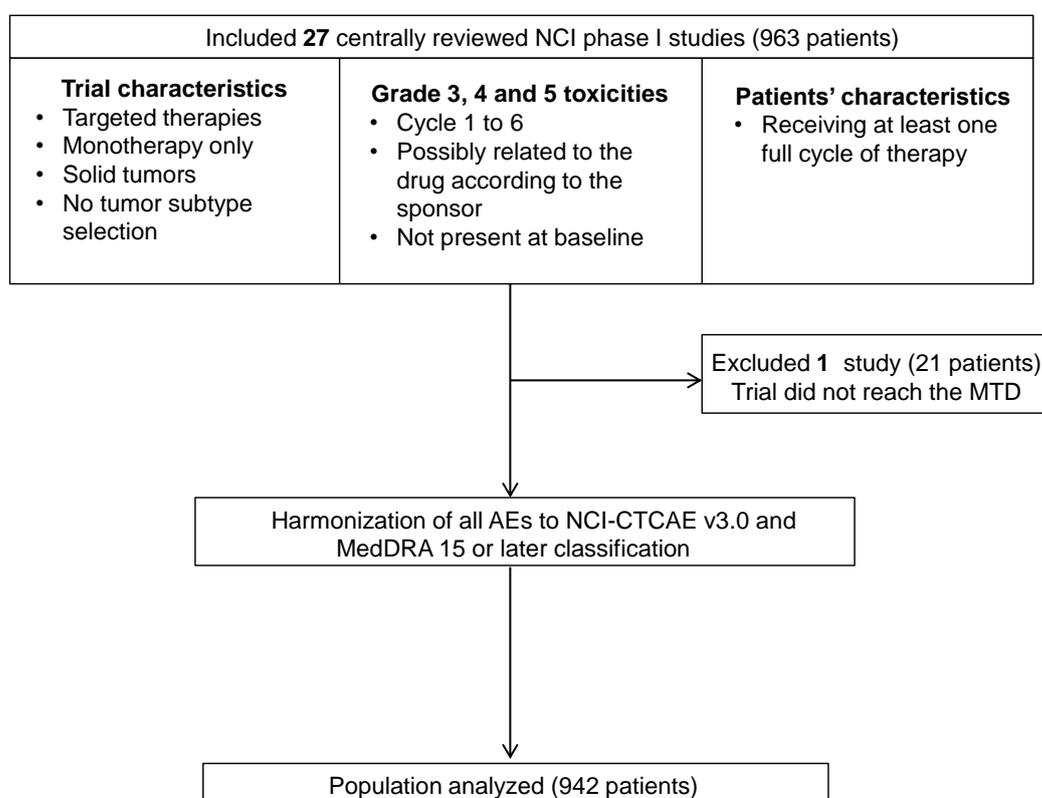


Figure 4.1: Flowchart of the study. Overview of the study design and trial selection. Abbreviations: AEs, Adverse Events; NCI, National Cancer Institute; NCICTCAE, NCI Common Terminology Criteria for Adverse Events; MedDRA, Medical Dictionary for Regulatory Activities.

ders, 4% dermatological, 3% liver, 3% glycemia, 3% vascular disorders and others. More specifically, in cycle 1, 15.1% of patients had a first non-hematologic toxicity, 6.7% a first hematologic toxicity, and 3.5% had both of them as compared to cycle 6 during which 1.5% first toxicities of each type were recorded, among patients still at risk (Table C.2).

Table 4.1: Trial characteristics

Characteristics	N (%)
Agent class	
Antiangiogenic agent	2 (7.7)
Antivascular agent	1 (3.8)
CDK inhibitor	2 (7.7)
HDAC inhibitor	3 (11.5)
HSP inhibitor	6 (23.1)
Immunotherapy	3 (11.5)
Monoclonal antibody	3 (11.5)
Proteasome inhibitor	2 (7.7)
Other	4 (15.4)
Administration route	
Intravenous	21 (80.8)
Oral	4 (15.4)
Intraperitoneal	1 (3.9)
Cycle duration	
14 days	2 (7.7)
21 days	6 (23.1)
28 days	15 (57.7)
42 days	3 (11.5)

Of the 942 patients treated in the 1st cycle, 97 had a DLT. Of them, 22 had been allocated to the MTD, 17 to a dose below and 58 to a dose above the MTD.

Table 4.2: Number (#) and percentage (%) of patients at each treatment cycle.

	cycle 1	cycle 2	cycle 3	cycle 4	cycle5	cycle 6
# patients	942	548	236	155	93	68
% patients	100	58.2	25.1	16.5	9.9	7.2

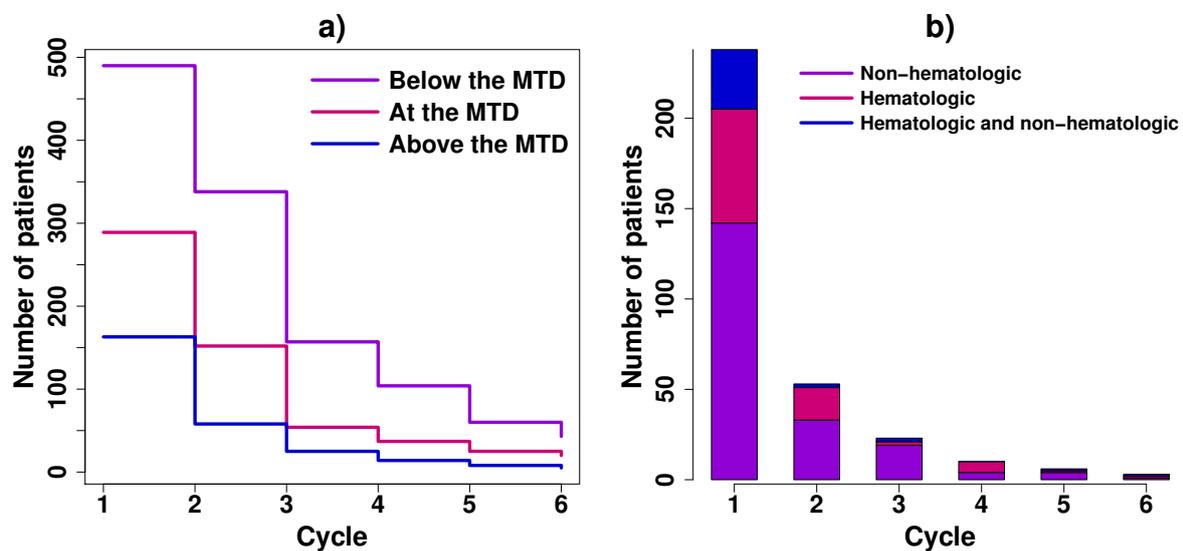


Figure 4.2: a) Patients per treatment cycle for the dose subgroups; below, above and at the MTD. b) Patients that had a toxic event per treatment cycle.

4.4.4 Time-to-first severe toxicity

The per-cycle risk of first severe toxicity decreased with treatment cycle (Table C.3). For patients allocated to the MTD the probability of a severe toxicity at the first treatment cycle was 27.3% [95% prediction interval: 22.6%; 32.1%] and this number monotonically decreased from 18.1% [14.7%; 21.6%] at cycle 2 down to 1.6% [0.2%; 5.5%] at cycle 6 (Figure 4.3 a), Table C.4). Of note, the risk of DLT in the first cycle, was 7.6% for

patients treated at the MTD, 3.5% and 35.6% for patients treated at doses below and above the MTD, respectively. The CIF of severe toxicity for patients treated at the MTD increased from 27.3% [22.6%; 32.1%] at cycle 1 to 52.9% [43.7%; 61.5%] at cycle 6. For patients assigned to doses below the MTD, CIF increased from 12.1% [9.7%; 15.2%] to 33.3% [27%; 40.5%] and to doses above the MTD from 48.9% [40.4%; 56.7%] to 80.1% [70.1%; 88.7%] (Figure 4.3 b)).

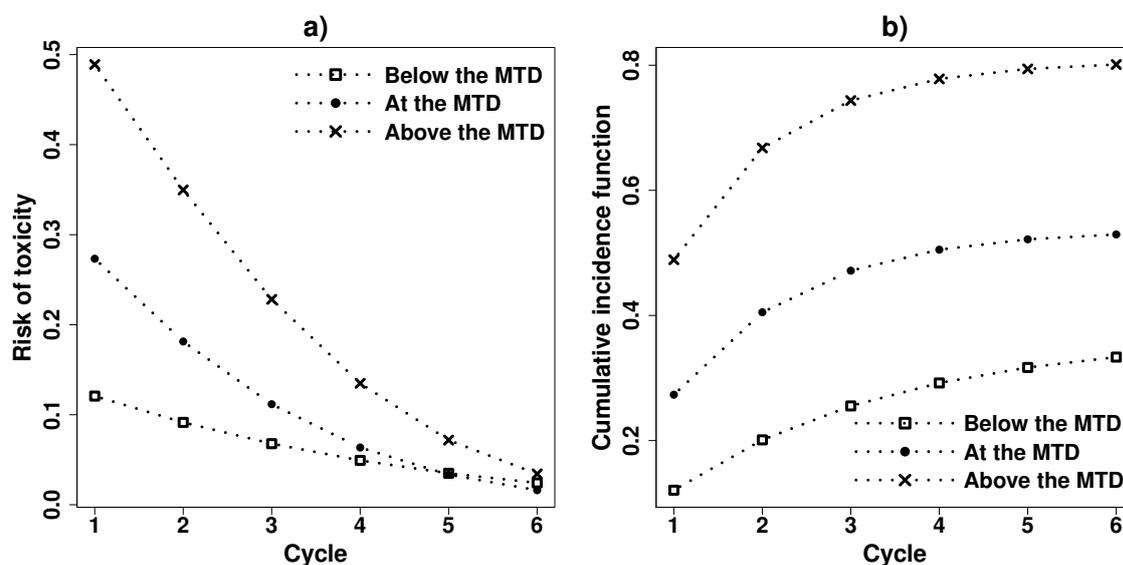


Figure 4.3: a) Risk and b) cumulative incidence function of severe toxicity over 6 cycles, for groups of patients treated at doses below, above and at the MTD.

Results of the complete model 4.1 are presented in Table C.5. Adjusting the model for both cycle and dose did yield similar results to those presented in the previous analysis. Figure 4.4 presents the risk and the CIF of severe toxicity for patients treated at the MTD for the 2 models. We can see that the predicted CIF of toxicity was close to the observed one. Thus, we conclude to small bias concerning the results of the complete model. Per-cycle results for patients at the MTD are given in Table C.6. The nomogram obtained from the complete model relating the CIF at cycle 6 to the risk at cycle 1 is given in Table 4.3. It shows that for a risk of severe toxicity of 5% in the first cycle of

treatment, the predicted CIF over 6 cycles of treatment was 10.7% [7.1%; 15.8%]. Most importantly, for the traditionally accepted thresholds of 20% and 30% of severe toxicity in cycle 1, targeted to identify the RP2D, the predicted CIFs at cycle 6 were 44.1% [39.3%; 49.5%] and 63% [56.8%; 69.5%], respectively.

Results of the residual check for the cycle and dose linearity can be found in Tables C.7 - C.10 of Appendix C. The per-cycle risk of toxicity and the CIF, among the different cycle effects, is presented in Table C.8. Assuming a different cycle effect for cycles 1-3 and 4-6 does not yield different results from the model, where we consider time to be linear on the probit scale. Assuming a different cycle effect for cycles 1-2 and 3-6 gives slightly higher rates of toxicity. Assuming a different dose effect below the MTD and at and above the MTD give very similar results to the case where we consider dose to be linear on the probit

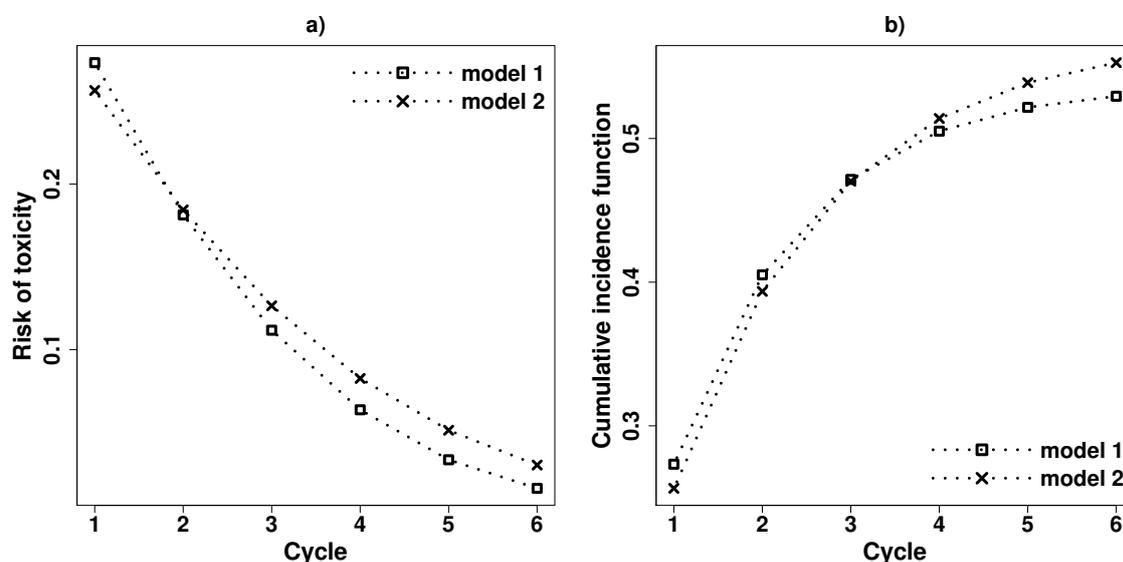


Figure 4.4: Risk and cumulative incidence function of severe toxicity at the MTD. Model 1 refers to the dose subgroup model and Model 2 to the complete model.

Next, we fitted the complete model 4.1 for 17 trials separately, for which we had enough data. We estimated the risk and the CIF of severe toxicity for patients allocated

Table 4.3: Cumulative incidence function of severe toxicity, assuming that risk of severe toxicity in the first cycle ranges between 5% and 35%.

cycle 1	cycle 2	cycle 3	cycle 4	cycle 5	cycle 6
0.050	0.078	0.093	0.101	0.105	0.107
0.100	0.157	0.189	0.207	0.217	0.221
0.150	0.235	0.284	0.311	0.326	0.334
0.200	0.311	0.374	0.410	0.430	0.441
0.250	0.384	0.460	0.503	0.527	0.541
0.300	0.455	0.539	0.587	0.615	0.630
0.350	0.522	0.613	0.664	0.693	0.710

to the MTD (Figure 4.5 - 4.6). After 6 cycles of treatment, among the 17 trials, the CIF exceeded 50% in 12 of them and remained below 30% in only 1.

4.4.5 CIF of hematologic and non-hematologic severe toxicity

Results of the “naive” approach comparing non-hematologic and hematologic severe toxicity, assuming each time that one of the two was not of interest can be found in Table C.11. For patients at the MTD, the CIF of non-hematologic severe toxicity by the end of cycle 6 was 48.3% [38.9%; 59.0%], twice as high as that of hematologic toxicity 24.3% [17.8%; 33.1%] (Figure 4.7 b), Table C.12). Similar results were observed for the other two dose subgroups. For patients below the MTD, the CIF of non-hematologic toxicity after 6 cycles reached 23.8% [18.2%; 30.7%], whereas the CIF of hematologic toxicity was 14.8% [9.9%; 20.7%] (Figure 4.7 a)). Finally, for patients above the MTD, the CIF of non-hematologic toxicity after 6 cycles was 72% [61.2%; 82.2%] versus 39.6% [27.7%; 53.6%] for hematologic toxicity (Figure 4.7 c)).

Next we proceeded with the competing risks analysis, estimating the cause-specific hazards of strictly non-hematologic and hematologic severe toxicity (Table C.13). For

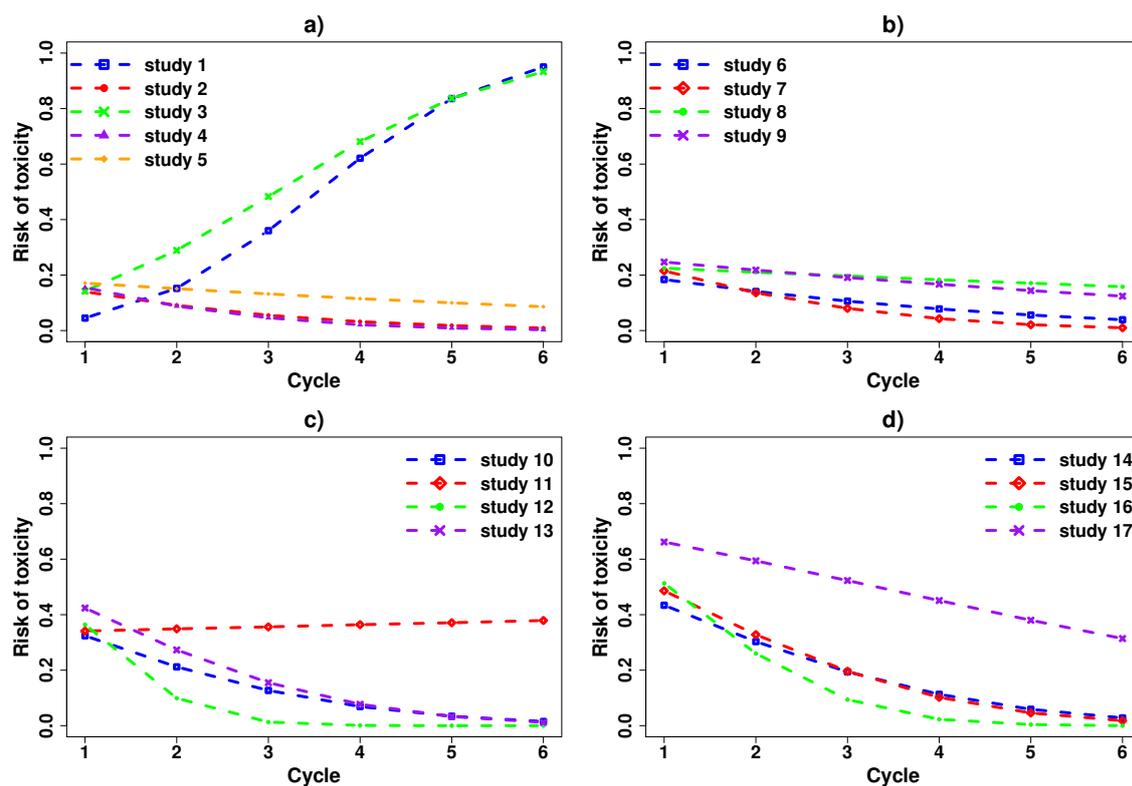


Figure 4.5: Risk of severe toxicity at the MTD per study.

patients assigned to the MTD level, the CIF of non-hematologic severe toxicity by the end of cycle 6 was 34.8% [26.6%; 44.1%] and was almost twice as much as that of having hematologic severe toxicity (18.2% [12.8%; 23.9%]) (Figure 4.8 b), Table C.14). Similar results were observed in patients treated below and above the MTD. In particular, for patients below the MTD, the CIF of exclusively non-hematologic severe toxicity reached 20.2% [15.2%; 26.1%] versus 13.1% [8.8%; 18.7%] for hematologic toxicity. For patients above the MTD it reached 45.6% [36.4%; 55.1%] for exclusively non-hematologic severe toxicity and 34.2% [24.8%; 44.5%] for hematologic toxicity (Figure 4.8 a), c).

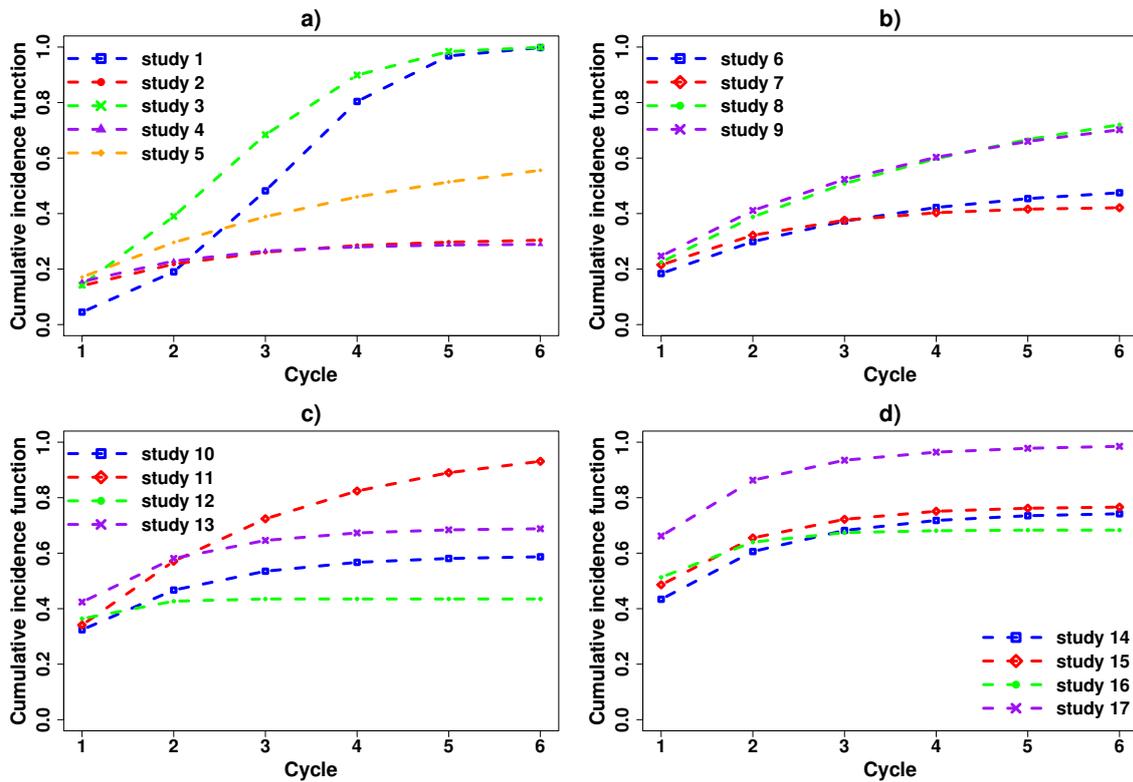


Figure 4.6: Cumulative incidence function of severe toxicity at the MTD per study.

4.4.6 Sensitivity analysis

Reclassification of mixed cases into non-hematologic toxicities led, as expected, to inflation of the CIF, of competing risks, of non-hematologic toxicity (Table C.15 for cause-specific hazards). At the MTD, it reached 39.2% [31.2%; 48.9%] and 13.8% [9.2%; 19.6%], for non-hematologic and hematologic respectively (Figure 4.9 b), Table C.16). For patients allocated to doses below the MTD, the CIF of non-hematologic toxicity by cycle 6 was 21.8% [18.1%; 29%] versus 11.6% [7.5%; 16.5%] for hematologic toxicity (Figure 4.9 a)). Above the MTD, the CIF of non-hematologic toxicity reached 63.2% [52.4%; 73.2%], whereas that of hematologic was 17.1% [10.3%; 25.5%] (Figure 4.9 c)).

For the second sensitivity analysis, exclusion of hematologic grade 3 AEs from the definition of severe toxicity led to a decrease in the number of hematologic severe toxicities and a small increase in that of non-hematologic. More precisely, in cycle 1, 17.5%

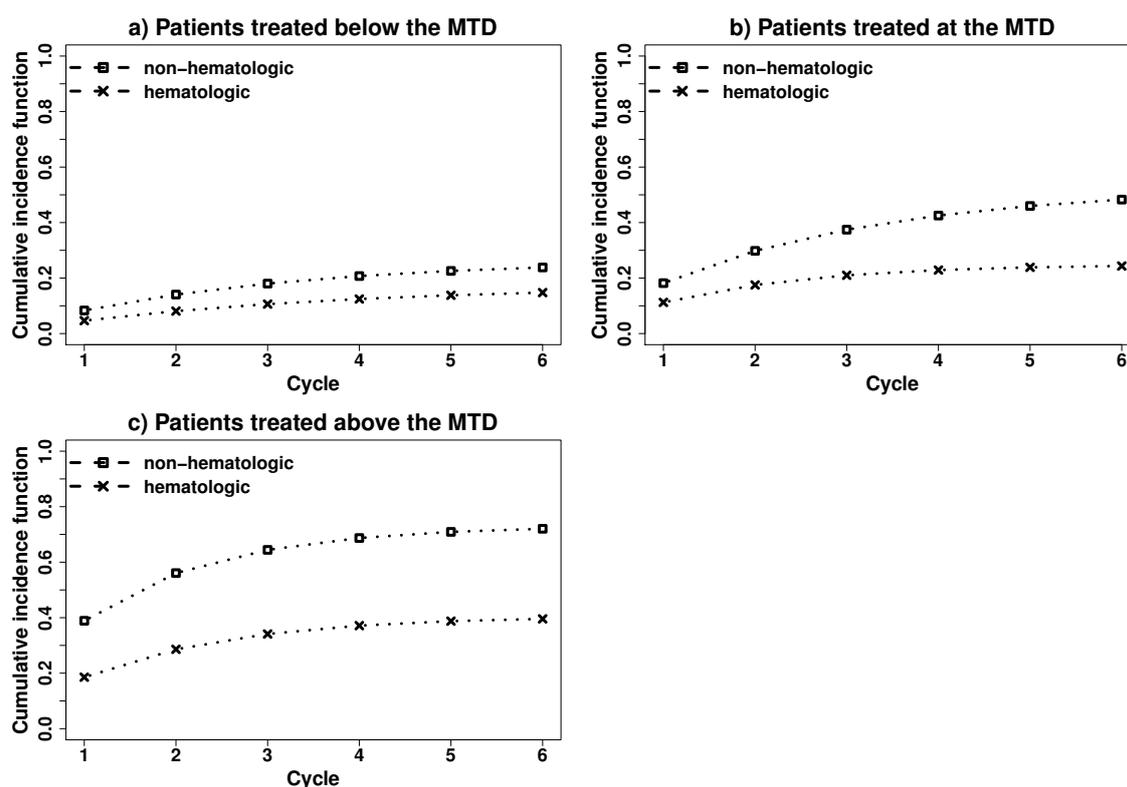


Figure 4.7: Cumulative incidence function for patients treated at doses below, above and at the MTD. CIFs were calculated for non-hematologic toxicity alone and hematologic toxicity alone.

of patients had a first non-hematologic toxicity (3.8% in cycle 6), 1% a first hematologic toxicity (0% in cycle 6), and 1.3% had both of them (1.3% in cycle 6) (Figure 4.10, Table C.17). After repeating the analysis for the 3 dose subgroups, separately (Table C.18), we found similar CIFs of any type of severe toxicity (Figure 4.11, Table C.19), as in our first analysis. However, the CIFs of strictly non-hematologic and hematologic toxicity were higher and lower respectively (Tables C.20 - C.21). At cycle 6, the ratio of the two CIFs was equal to 6. For instance, at the MTD, the CIF of strictly non-hematologic toxicity by cycle 6 reached 45% [35.6%; 55.0%] versus 7.5% [3.2%; 14.0%] for hematologic toxicity (Figure 4.12 b)). Likewise, below and above the MTD, the CIFs of strictly non-hematologic and hematologic toxicity, by cycle 6, were 22.8% [17.3%; 29.5%] versus 3.3% [1.2%; 7.5%] and 61.4% [52.1%; 71.5%] versus 13.6% [6.4%; 22.8%], respectively

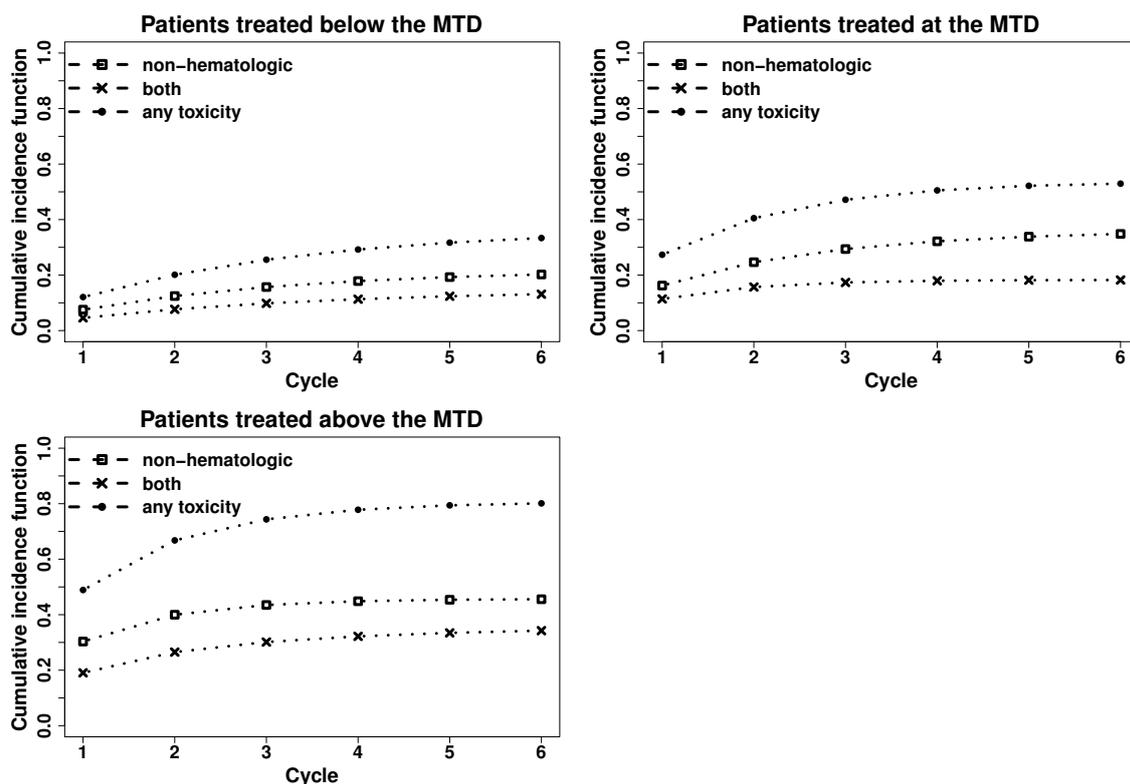


Figure 4.8: Cumulative incidence function for patients treated at doses below, above and at the MTD. CIFs were calculated for overall toxicity, non-hematologic toxicity and both hematologic and non-hematologic toxicity.

(Figure 4.12 a) , c)).

4.5 Discussion

In this chapter we have investigated the association of time-on-treatment, i.e. the number of treatment cycle, with the risk of severe toxicity. We showed that for patients assigned at the MTD, the risk of severe toxicity was 27% in cycle 1 and then per-cycle risk decreased for each successive cycle. This is in line with results from other studies (Postel-Vinay et al., 2011). At the MTD the CIF of severe toxicity by the end of cycle 6 was 53%. This risk is much higher than the 20%-33% risk of severe toxicity usually

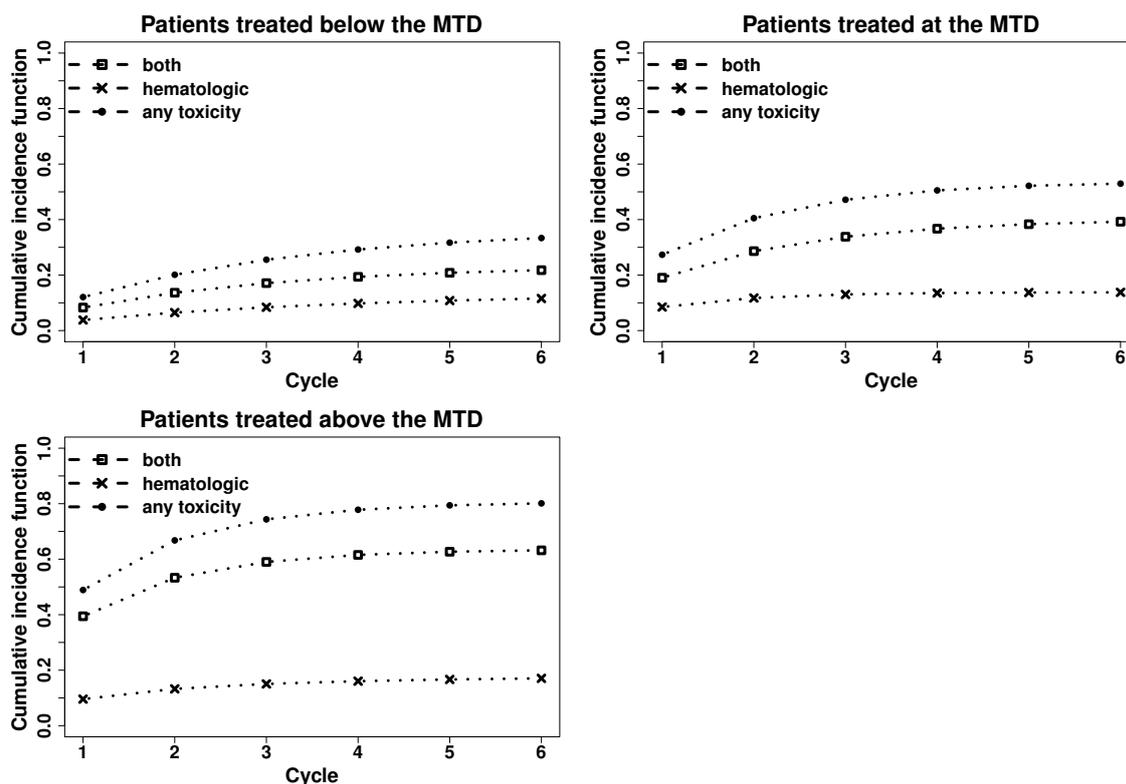


Figure 4.9: Cumulative incidence function for patients treated at doses below, above and at the MTD. CIFs were calculated for overall toxicity, both hematologic and non-hematologic toxicity and hematologic toxicity alone.

targeted in cycle 1 for the determination of the MTD and the RP2D. This CIF reached as much as 80% for doses above the MTD. We have shown the importance of considering separately the CIF of non-hematologic and hematologic severe toxicity, as the first one was always higher. What is more, at the MTD, the CIF of severe toxicity was 48% for non-hematologic toxicity, in the presence or not of hematologic toxicity and 24% for hematologic toxicity, in the presence or not of non-hematologic toxicity. At the MTD, the CIF of severe toxicity was made of 35% exclusively non-hematologic toxicity and of 18% hematologic or mixed toxicity.

The appropriate design and approaches to conduct phase I trials of MTAs and more recently, of immunotherapies is much debated. Many authors have underlined that delayed and cumulative toxicities of MTAs, resulting from the prolonged administration of

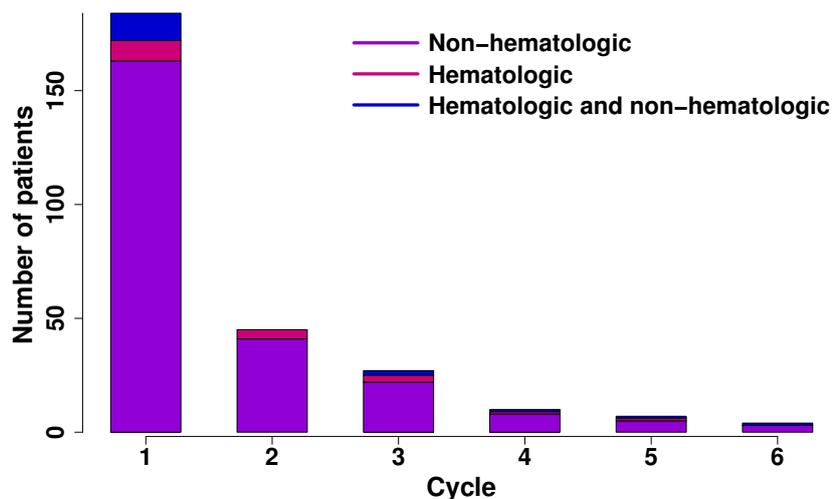


Figure 4.10: Patients that had a toxic event per treatment cycle.

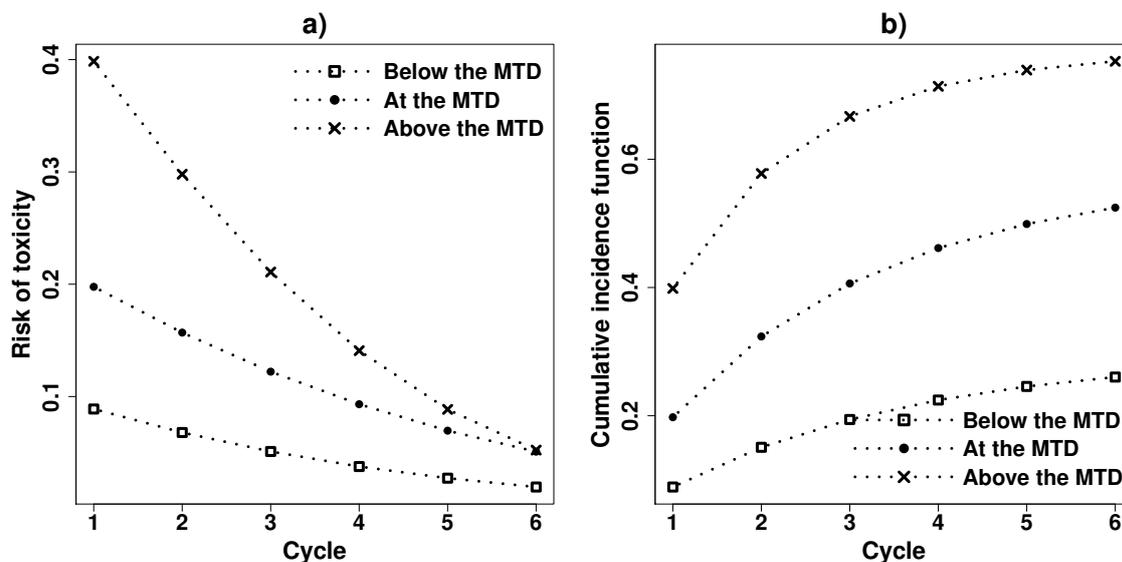


Figure 4.11: Risk and cumulative incidence function of severe toxicity over 6 cycles, for groups of patients treated at doses below, above and at the MTD.

these treatments, have a non-negligible impact on the MTD definition and on the selection of the RP2D. This raises the question of what is an adequate DLT period (Booth et al., 2008; Le Tourneau et al., 2010a; Soria, 2011). Illustrative of this is the poor pre-

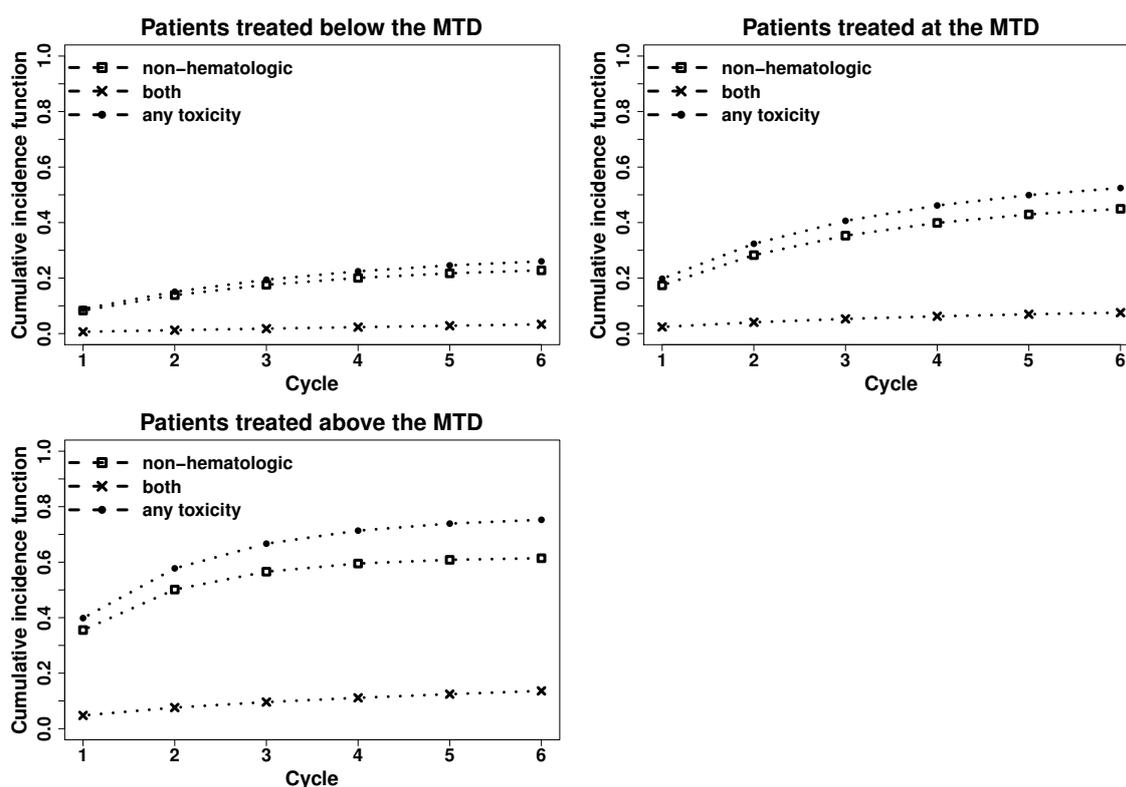


Figure 4.12: Cumulative incidence function of severe toxicity for patients treated at doses below, above and at the MTD. CIFs were calculated for overall toxicity, non-hematologic toxicity and both hematologic and non-hematologic toxicity.

diction of future approved dose levels from phase I and the resulting re-evaluation of the MTD in subsequent phases of treatment development (Le Tourneau et al., 2010b; Iasonos et al., 2012; Jardim et al., 2014).

The strengths of our study, besides the large number of trials and patients included, is the detailed information collected about e.g. planned doses, treatment cycles, grades and types of toxicity, etc.

Nevertheless, the fact that the treatments evaluated in the included phase I studies are not from the most recent classes of agents with late onset toxicity, such as immunotherapy, is a limitation of our study. Furthermore, we used a definition of severe toxicity that is probably much larger than the usual definition of the dose limiting toxicity that may exclude some grade 3 events such febrile neutropenia that would last less than 7 days.

Hence, it is likely that we have overestimated the CIF of severe toxicity. Due to the fact that DLTs are generally not recorded after cycle 1 and we did not have access to all protocols, we cannot predict what is the extent of this overestimation.

For phase I designs we believe that 3-6 cycles should be considered in order to more reliably determine the RP2D. Conversely, the interest of a treatment that would not result in disease stabilization after two cycles would be questionable. In our analysis the risk of grade 3 to 5 toxicity of 20% to 30% on the 1st treatment cycle translated into cumulative incidence function of such events of 44% to 63% at the end of the 6th treatment cycle. Therefore, a reasonable target of cumulative risk of severe toxicity may be of the order of 40% to 45% over 6 cycles. We then suggest that re-evaluation of the CIF of severe toxicity should be part of the stated objectives of the now rather popular expansion cohorts (Dahlberg et al., 2014). Several dose-escalation methods have been proposed that take into account time-on-treatment in the dose-escalation process, such as time-to-event CRM, the CRM for longitudinal data (Paoletti et al., 2015) and of course the designs presented in Chapters 2 and 3. Nevertheless, as thoroughly discussed before, toxicity is only one component of the overall evaluation leading to the definition of the optimal dose of a treatment. Pharmacokinetic data or biomarker measurements are also important factors that help refine the dose selected for the next phase of treatment development. Designs that make use of all collected data should improve the efficiency of phase I trials of MTAs at defining the RP2D and therefore avoid the need to re-evaluate accepted doses in subsequent phases.

Chapter 5

Conclusion

The objective of this thesis was to develop adaptive designs that can be applied in phase I/II trials of molecularly targeted agents in oncology. Phase I trials are the cornerstone of drug development and their correct design and conduct can help improve the overall procedure. A major step in order to improve phase I trials is to integrate all information collected during the trial. First of all, we were interested in developing a design, for the identification of the OD that can consider several treatment cycles and not just the first one, as it is currently done. The principal goal was to account for late-onset and cumulative toxicities. Furthermore, the extensive use of biomarkers that measure the drug activity, requires alternative approaches, which consider the biomarker as a continuous variable, measured at multiple time points during the course of the trial, so as to capture the information provided. It is well known that dichotomization of such endpoints can only lead to substantial loss of information. What is more, considering only a single biomarker measurement is not sufficient to define the drug activity, since the course of the biomarker changes in time, indicating first response to treatment and eventually disease progression. Another important issue that was of interest to that thesis was censored data. When the OD selection is defined based on a long period, i.e. 3-6 cycles, encountering missing data is inevitable.

For the OD selection we proposed an adaptive design that combines both toxicity and activity data, measured over a period of six treatment cycles. We implemented a recent joint modeling technique of time-to-DLT outcome and continuous repeated biomarker

measurements, under a shared random effect. This joint model estimation was using likelihood inference and model parameters were obtained from the numerical maximization of the likelihood. The MTD was defined cumulatively, over six cycles, and the OD was the lowest within a range of active and safe doses.

To investigate the robustness of the design and its ability to correctly identify the OD we assessed its operating characteristics in a wide range of scenarios. We considered smaller sample sizes, larger variances, altered random effects distribution, we assumed an increasing hazard per treatment cycle and most importantly we generated several scenarios under model misspecification. The method was highly efficacious regarding the correct OD selection. Even under excessive censoring, with only 7% of patients arriving at the end of the trial, the method was robust and estimations had small bias. What is more, our design was safe, since study participants were generally not exposed to highly toxic or subtherapeutic doses. This was the first design for phase I trials that included several treatment cycles, continuous and repeated biomarker measurements, a cumulative definition for the MTD and tackled the issue of missing responses.

A limitation of this design is associated to the dose-activity model. We implemented a linear mixed effects model with a log term for the dose so as to allow for a dose-activity relationship that was not strictly linear. Nevertheless, it was not enough to capture a plateau. Additionally, due to the large number of model parameters and the small sample size we encountered certain difficulties regarding the joint model convergence. This issue was apparent when the MTD was the first dose in the dose range.

To address the aforementioned limitations, we extended our adaptive design by allowing for model flexibility in the dose-activity relationship and a broader OD definition. The core of the design remained the same. To make the model suitable for an activity that plateaus after a certain dose level, we substituted the log model by a model that can capture both a plateau and a strictly linear dose-activity relationship.

Investigation of the robustness of the new design was focused on the same motivating example of the ovarian cancer trial so as to allow for comparison of the two approaches. Extensive simulations showed that the new model performed very well under various degrees of model misspecification. An important feature of the plateau model was that

when data was generated from the log model, the percentage of correct OD recommendation was almost identical to that observed when analyzing data with the log model. That was an indication of the model's adaptability to different activity biomarker trajectories. Last but not least, we addressed the case of convergence failure. Under the log model, certain simulations could not proceed with the joint modeling, whereas with the plateau model and for the same simulations no such issue was observed.

The proposed adaptive designs could be applied to different settings from the ones we investigated, such that of combinational therapies. Under this particular framework, we combine a new treatment with one or more of the standard ones. Our models could be adjusted to allow for more than one dose variable, as well as dose interactions in case of synergistic effects. However, increasing the number of variables makes parameter estimation more challenging, especially in the beginning of the trial. In that case, it could be useful to use joint modeling but with simpler models or to turn to the Bayesian inference, so as to incorporate prior information regarding the standard treatments. In addition to this, the adaptive designs could also be employed when we work with more than one measures of activity, i.e. multiple biomarkers. In that case there are two possible solutions. The first one is to create a score of all activity outcomes and to use that score as the outcome of the activity model. This is a quite common practice. The second option, if for example we have two different biomarkers, would be to consider separately both activity endpoints. In that case the joint models would consist of 3 models, two for the activity and one for the toxicity.

A limitation of both adaptive designs is related to their lack of comparability with existing methods. Several dose finding methods have been proposed in the literature. There exist methods that account for a longer DLT period, without however considering several treatment cycles or a cumulative definition for the MTD. There are also methods that consider activity as a continuous endpoint, but with a single measurement per patient. Therefore, comparing our design with alternative approaches was not feasible, since it would not be possible to create matching scenarios. Another limitation is with regard to the joint modeling method. We mentioned earlier that Barrett et al. (2015) used the probit model for toxicity and the linear mixed effects model for activity with the

purpose to exploit specific properties of the skew normal distribution. Their aim was to express the joint likelihood in a closed form, by integrating out the random effects. Although this property allowed for exact likelihood inference, joint modeling was restricted to shared random effects. Under shared random effects, we make the assumption that the same random effect accounts for both the association between the longitudinal and event outcomes, and the correlation between the repeated measurements. On the contrary, there exist joint modeling techniques that are more flexible (Rizopoulos, 2012), and can account for various kinds of shared quantities, such as shared random effects and shared model parameters. Nonetheless, these methods come with certain hindrances, which are either associated to computational difficulties or lack of interpretability of the results.

The majority of the existing dose finding methods, presented in Chapter 1, relies on the Bayesian framework. Bayesian inference is appropriate for parameter estimation, when prior information exists regarding the toxicity and activity endpoints. An example is phase I trials in pediatrics, where the drug under investigation has already been tested in adults. However, in most first in human clinical trials, there is little or no knowledge with respect to toxicity and activity. An impediment, when considering both outcomes, is that we need to elicit more priors than we normally would, if we considered only toxicity. Under no prior knowledge, prior distributions are most often non-informative. This is an issue under small sample sizes, because parameter estimation, in the absence of data, depends only on priors. This is why, certain dose-finding methods, such as Riviere et al. (2016) or Yuan and Yin (2009), start the trial with a rule-based design that imitates the ‘3+3’ or similar designs, with the purpose to obtain some data, before proceeding with the Bayesian inference. This is not very different from our design that starts with the ‘2+2’ design and when there is enough data it switches to the joint modeling. Nevertheless, an advantage of our design is that it does not rely on approximations to estimate the model parameters, as it is the case for Bayesian methods, but on numerical maximization.

For our design, we relied on jointly modeling toxicity and activity for the OD selection. It has been shown that when two outcomes are correlated parameter estimation can significantly benefit from the use of joint modeling. There have been several articles in the literature (Schluchter, 1992; Faucett and Thomas, 1996; Wulfsohn and Tsiatis, 1997;

Henderson et al., 2000; Tsiatis and Davidian, 2004; Rizopoulos, 2012) showing that joint models provide more robust estimates, compared to models that consider the outcomes independently. Joint modeling is a suitable way to correct the toxicity model for the missing data, observed due to disease progression. However, a question raised here is whether joint models are absolutely necessary for dose selection. For instance, when we target 40% of cumulative toxicity, we aim to find a dose with a probability of DLT as close as possible to the target. It is possible that even if estimates are obtained from independent models, the recommended dose will still be the same. Nonetheless, we showed in a simulation study that most often under correlated outcomes joint modeling outperforms independent modeling, despite the estimation of the extra parameter. Depending on the data variability, it is possible to observe differences between the two approaches even in the presence of a weak correlation. It would be interesting though to further investigate the value of joint modeling in phase I/II clinical trials, within a broader range of scenarios. In the case where the two outcomes are not correlated considering them separately could be a more appropriate approach.

Another point that drew our attention was model predictions. When we estimated the MTD and OD, we used a conditional or else subject-specific model for both activity and toxicity models. When applying a conditional model we make the assumption that the patient entering the study is most likely to have a random effect that approximates a mean value of zero. An alternative approach would be to estimate the marginal or else the population-averaged expectation, by integrating over the random effect distribution. Due to the symmetry of the normal distribution, both approaches, conditional and marginal, give the same predictions for the linear outcome. However, this is not the case for the survival outcome. Thus, the first question raised is which of the two approaches would be more suitable for the survival outcome. We chose the conditional expectation, since we wanted to select a dose for the next patient and not generalize the results on the population. Another matter worth mentioning is whether we should focus on an immortal cohort or on partly conditional predictions. We remind that an immortal cohort implies that subjects under study would not drop out due to toxicity or lack of activity. This scenario would be more realistic in a phase II or a phase III clinical trial, where

patients would probably receive co-medication in the presence of serious adverse events. On the contrary, partly conditional effects, population-averaged or subject-specific, are estimated conditionally on the fact that subjects are still in the study. For example, a partly conditional population-averaged expectation would be appropriate if we wanted to predict the risk of toxicity in a specific treatment cycle. We mentioned above that for a particular patient entering the study we used the subject-specific expectation. However, if we wished to make individual predictions for patients already enrolled in the study we would choose the partly conditional subject-specific expectation. This is because we have already followed the patient and there exists information regarding his random effect. Using joint modeling, we can predict at the same time the toxicity and activity outcome of a patient, so as to allow for dose modification during the course of the trial, or otherwise inpatient dose modification. That would be very important, given the fact that patients can be treated for several treatment cycles. We could potentially identify patients for whom the drug can be proved to be highly toxic and therefore, decrease the dose so as to avoid future dose limiting toxicities. In a similar way, if a patient does not respond to the treatment and a higher dose is deemed to be safe, then it would be possible to allow for dose escalation for that particular patient. As a result, patients could remain for a longer period in the trial and potentially benefit from the treatment.

The last part of this thesis focused on the analysis of 27 completed phase I trials of targeted therapies as monotherapy. These trials were conducted by the NCI and were provided to the EORTC and the DLT-TARGETT. In the previous chapters, we discussed the importance of considering time when conducting a phase I trial. However, till now there was nothing in the literature, documenting the risk of severe toxicity, for more than one treatment cycles. Therefore, our primary interest was to estimate the per-cycle risk, as well as the cumulative incidence function of severe toxicity over up to six treatment cycles. A secondary objective was to estimate the same quantities, but separately for hematologic and non-hematologic severe toxicities. These types of toxicities do not occur in the same way and do not have the same impact on the human body, thus there is merit in estimating the risk separately.

To that end, we implemented a probit time-to-toxicity model that can estimate both

the per-cycle risk and the cumulative incidence function. We performed the analysis, first separately for groups of patients treated at doses below, above, and at the MTD and then overall patients and we provided a nomogram. We investigated the risk of toxicity for hematologic and non-hematologic toxicities, and finally, we conducted a sensitivity analysis. We showed that the cumulative incidence function of severe toxicity for patients treated at the MTD was 27% at the end of the first cycle and reached 53% by the end of cycle six. For patients administered to the MTD, the cumulative incidence function of non-hematologic and hematologic toxicity was 35% and 18%, accordingly.

This project was the first one to provide results on the per-cycle risk and the cumulative incidence function of severe toxicity over several treatment cycles. We additionally, provided guidelines for the good conduct of phase I trials, and reasonable target toxicity levels, based on the results from the DLT-TARGETT. A limitation of this analysis is related to the severe toxicity definition. Traditionally, the MTD is defined based on the DLTs occurring during the first treatment cycle. DLTs are defined upon on certain grade 3 and grade 4 toxicities and they are study specific. We were interested in the risk of toxicity over several cycles, and for that we would need to reconstruct the DLT variable for all six cycles. However, we did not have access to all 27 study protocols. As a consequence, for the analysis we focused on the first severe toxicity, which does not necessarily coincide with a DLT. Therefore, considering the first severe toxicity could have resulted in an overestimation of the risk of toxicity. What is more, information recorded provided the dose limiting toxicities that occurred during the first cycle and for the following cycles patients were censored. Hence, we could not know if patients were removed from the study due to excessive toxicity, disease progression or consent withdrawal.

It has been shown that the MTD of targeted therapies, contrary to cytotoxic therapies, is quite often re-evaluated in subsequent phases (Le Tourneau et al., 2010b; Iasonos et al., 2012; Jardim et al., 2014). This is a major problem of phase I trials of these agents, indicating that new methods and new designs should be applied when recommending the MTD and the OD. With this thesis, we wished to provide evidence that there is space for improvement and to that end we developed an adaptive design and an extension of that design that could be applied to these trials. Finally, after the analysis of the DLT-

TARGETT, it became apparent that for the definition of the recommended phase II dose, it is vital to consider 3-6 treatment cycles, consider a cumulative definition of the target toxicity level and if possible, incorporate activity measurements. In conclusion, integration of more information could lead to more efficient phase I trials of targeted therapies.

Bibliography

- Abola, M. V., Prasad, V., and Jena, A. B. (2014). Association between treatment toxicity and outcomes in oncology clinical trials. *Annals of Oncology*, 25(11):2284–2289.
- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. *International Symposium on Information Theory*, pages 267–281.
- Arnold, B. C. (2009). Flexible univariate and multivariate models based on hidden truncation. *Journal of Statistical Planning and Inference*, 139(11):3741–3749.
- Azzalini, A. (2005). The skew-normal distribution and related multivariate families (with discussion). *Scandinavian Journal of Statistics*, 32:159–200.
- Babb, J., Rogatko, A., and Zacks, S. (1998). Cancer phase i clinical trials: efficient dose escalation with overdose control. *Stat Med*, 17(10):1103–1120.
- Barrett, J., Diggle, P., Henderson, R., and Taylor-Robinson, D. (2015). Joint modelling of repeated measurements and time-to-event outcomes: Flexible model specification and exact likelihood inference. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 77(1):131–148.
- Bekele, B. N. and Shen, Y. (2005). A Bayesian approach to jointly modeling toxicity and biomarker expression in a phase I/II dose-finding trial. *Biometrics*, 61(2):344–354.
- Berry, S. M., Carlin, B. P., Lee, J. J., and Muller, P. (2010). *Bayesian adaptive methods for clinical trials*. CRC Press.
- Booth, C. M., Calvert, A. H., Giaccone, G., Lobbezoo, M. W., Seymour, L. K., and Eisenhauer, E. A. (2008). Endpoints and other considerations in phase I studies of

- targeted anticancer therapy: Recommendations from the task force on Methodology for the Development of Innovative Cancer Therapies (MDICT). *European Journal of Cancer*, 44(1):19–24.
- Cai, C., Yuan, Y., and Ji, Y. (2014). A Bayesian Dose-finding Design for Oncology Clinical Trials of Combinational Biological Agents. *Journal of the Royal Statistical Society. Series C, Applied statistics*, 63(1):159–173.
- Champiat, S., Lambotte, O., Barreau, E., Belkhir, R., Berdelou, A., Carbonnel, F., Cauquil, C., Chanson, P., Collins, M., Durrbach, A., Ederhy, S., Feuillet, S., François, H., Lazarovici, J., Le Pavec, J., De Martin, E., Mateus, C., Michot, J. M., Samuel, D., Soria, J. C., Robert, C., Eggermont, A., and Marabelle, A. (2016). Management of immune checkpoint blockade dysimmune toxicities: A collaborative position paper. *Annals of Oncology*, 27(4):559–574.
- Cheng, L., Lopez-Beltran, A., Massari, F., MacLennan, G. T., and Montironi, R. (2018). Molecular testing for BRAF mutations to inform melanoma treatment decisions: A move toward precision medicine. *Modern Pathology*, 31(1):24–38.
- Cheung, Y. K. (2011). *Dose finding by the continual reassessment method*. Biostatistics Series, Chapman & Hall/CRC.
- Cheung, Y. K. and Chappell, R. (2000). Sequential designs for phase i clinical trials with late-onset toxicities. *Biometrics*, 56(4):1177–1182.
- Cheung, Y. K. and Chappell, R. (2002). A simple technique to evaluate model sensitivity in the continual reassessment method. *Biometrics*, 58:671–674.
- Crowley, J. and Hoering, A. (2012). *Handbook of statistics in clinical oncology*. Third Edition, Chapman & Hall/CRC.
- Cunanan, K. and Koopmeiners, J. S. (2014). Evaluating the performance of copula models in phase I-II clinical trials under model misspecification. *BMC Medical Research Methodology*, 14:1471–2288.

- Dahlberg, S. E., Shapiro, G. I., Clark, J. W., and Johnson, B. E. (2014). Evaluation of statistical designs in phase i expansion cohorts: The Dana-Farber/Harvard cancer center experience. *Journal of the National Cancer Institute*, 106(7).
- Doussau, A., Asselain, B., Le Deley, M. C., Geoerger, B., Doz, F., Vassal, G., and Paoletti, X. (2012). Dose-finding designs in pediatric phase I clinical trials: Comparison by simulations in a realistic timeline framework. *Contemporary Clinical Trials*, 33(4):657–665.
- Efron, B. (1981). Nonparametric estimates of standard error: The jackknife, the bootstrap and other methods. *Biometrika*, 68(3):589–599.
- Eisenhauer, E. A., Twelves, C., and Buyse, M. (2015). *Phase I cancer clinical trials: A practical guide*. Oxford University Press.
- European Medicines Agency (2016). Draft Guideline on the evaluation of anticancer medicinal products in man. Technical report.
- Faucett, C. L. and Thomas, D. C. (1996). Simultaneously modelling censored survival data and repeatedly measured covariates: A Gibbs sampling approach. *Statistics in Medicine*, 15(15):1663–1685.
- Henderson, R., Diggle, P., and Dobson, A. (2000). Joint modelling of longitudinal measurements and event time data. *Biostatistics*, 1:465–480.
- Hirakawa, A. (2012). An adaptive dose-finding approach for correlated bivariate binary and continuous outcomes in phase I oncology trials. *Statistics in Medicine*, 31(6):516–532.
- Hunsberger, S., Rubinstein, L. V., Dancey, J., and Korn, E. L. (2005). Dose escalation trial designs based on a molecularly targeted endpoint. *Statistics in Medicine*, 24(14):2171–2181.
- Iasonos, A., Gounder, M., Spriggs, D. R., Gerecitano, J. F., Hyman, D. M., Zohar, S., and O’Quigley, J. (2012). The impact of non-drug-related toxicities on the estimation of the maximum tolerated dose in phase i trials. *Clinical Cancer Research*, 18(19):5179–5187.

- Ibrahim, J. G., Chu, H., and Chen, L. M. (2010). Basic concepts and methods for joint models of longitudinal and survival data. *Journal of Clinical Oncology*, 28(16):2796–2801.
- ICH (1994). *Clinical safety data management: Definitions and standards for expedited reporting E2A*. <https://www.imim.es/media/upload/arxius/MEDIA436.pdf>. [Accessed: 20/09/2016].
- Ivanova, A. and Kim, S. H. (2009). Dose finding for continuous and ordinal outcomes with a monotone objective function: A unified approach. *Biometrics*, 65(1):1–19.
- Ivanova, A., Qaqish, B. F., and Schell, M. J. (2005). Continuous toxicity monitoring in phase II trials in oncology. *Biometrics*, 61(2):540–545.
- Ivanova, A. and Xiao, C. (2013). Dose-finding when the target dose is on a plateau of a dose-response curve: Comparison of fully sequential designs. *Pharmaceutical Statistics*, 12(5):309–314.
- James, K., Eisenhauer, E., Christian, M., Terenziani, M., Vena, D., Muldal, A., and Therasse, P. (1999). Measuring response in solid tumors: Unidimensional versus bidimensional measurement. *Journal of the National Cancer Institute*, 91(6):523–528.
- Jardim, D. L., Hess, K. R., LoRusso, P., Kurzrock, R., and Hong, D. S. (2014). Predictive value of phase I trials for safety in later trials and final approved dose: Analysis of 61 approved cancer drugs. *Clinical Cancer Research*, 20(2):281–288.
- Ji, Y. and Wang, S. J. (2013). Modified toxicity probability interval design: a safer and more reliable method than the 3 + 3 design for practical phase I trials. *J Clin Oncol.*, 31(14):1785–1791.
- Jin, I. H., Liu, S., Thall, P. F., and Yuan, Y. (2014). Using data augmentation to facilitate conduct of phase I–II clinical trials with delayed outcomes. *Journal of the American Statistical Association*, 109(506):525–536.
- Koopmeiners, J. S. and Modiano, J. (2014). A Bayesian adaptive phase I–II clinical trial for evaluating efficacy and toxicity with delayed outcomes. *Clinical Trials*, 11(1):38–48.

- Kummar, S., Gutierrez, M., James, H., and Murgo, A. J. (2006). Drug development in oncology : classical cytotoxics and molecularly targeted agents. 62(1):15–26.
- Lawrence Gould, A., Boye, M. E., Crowther, M. J., Ibrahim, J. G., Quartey, G., Micallef, S., and Bois, F. Y. (2015). Joint modeling of survival and longitudinal non-survival data: Current methods and issues. Report of the DIA Bayesian joint modeling working group. *Statistics in Medicine*, 34(14):2181–2195.
- Le Tourneau, C., Diéras, V., Tresca, P., Cacheux, W., and Paoletti, X. (2010a). Current challenges for the early clinical development of anticancer drugs in the era of molecularly targeted agents. *Targeted Oncology*, 5(1):65–72.
- Le Tourneau, C., Stathis, A., Vidal, L., Moore, M. J., and Siu, L. L. (2010b). Choice of starting dose for molecularly targeted agents evaluated in first-in-human phase I cancer clinical trials. *Journal of Clinical Oncology*, 28(8):1401–1407.
- Lee, M., Feuer, E. J., and Fine, J. P. (2018). On the analysis of discrete time competing risks data. *Biometrics*.
- Luke, J. J. and Hodi, F. S. (2012). Vemurafenib and BRAF inhibition: A new class of treatment for metastatic melanoma. *Clinical Cancer Research*, 18(1):9–14.
- Mccrink, L. M., Marshall, A. H., and Cairns, K. J. (2013). Advances in joint modelling: A review of recent developments with application to the survival of end stage renal disease patients. *International Statistical Review*, 81(2):249–269.
- Mcculloch, C. E., Neuhaus, J. M., Olin, R. L., San, S. F., Francisco, S., and Francisco, S. (2016). Informative Visit Processes. 72(4):1315–1324.
- Mick, R. and Ratain, M. J. (1993). Model-guided determination of maximum tolerated dose in Phase I clinical trials: Evidence for increased precision. *J Nat Can Inst*, 85(3):217–223.
- National Cancer Institute (2018). *Targeted Cancer Therapies*. <https://www.cancer.gov/about-cancer/treatment/types/targeted-therapies/targeted-therapies-fact-sheet#q1> [Accessed: 20/05/2018].

- Neuenschwander, B., Branson, M., and Gsponer, T. (2008). Critical aspects of the bayesian approach to phase i cancer trials. *Stat Med*, 27(13):2420–2439.
- O’Quigley, J. and Chevret, S. (1991). Methods for dose finding studies in cancer clinical trials: A review and results of a monte carlo study. *Statistics in Medicine*, 10:1647–1664.
- O’Quigley, J., Hughes, M. D., and Fenton, T. (2001). Dose-finding designs for HIV studies. *Biometrics*, 57(4):1018–1029.
- O’Quigley, J., Pepe, M., and Fisher, L. (1990). Continual reassessment method: A practical design for phase 1 clinical trials in cancer. *Biometrics*, 46(1):33–48.
- O’Quigley, J. and Shen, L. Z. (1996). Continual reassessment method: A likelihood approach. *Biometrics*, 52(2):673–684.
- Paoletti, X., Doussau, A., Ezzalfani, M., Rizzo, E., and Thiébaud, R. (2015). Dose finding with longitudinal data: Simpler models, richer outcomes. *Statistics in Medicine*, 34(22):2983–2998.
- Paoletti, X. and Kramar, A. (2009). A comparison of model choices for the continual reassessment method in phase i cancer trials. *Statistics in Medicine*, 28:3012–3028.
- Postel-Vinay, S., Collette, L., Paoletti, X., Rizzo, E., Massard, C., Olmos, D., Fowst, C., Levy, B., Mancini, P., Lacombe, D., Ivy, P., Seymour, L., Le Tourneau, C., Siu, L. L., Kaye, S. B., Verweij, J., and Soria, J. C. (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. *European Journal of Cancer*, 50(12):2040–2049.
- Postel-Vinay, S., Gomez-Roca, C., Molife, L. R., Anghan, B., Levy, A., Judson, I., De Bono, J., Soria, J. C., Kaye, S., and Paoletti, X. (2011). Phase I trials of molecularly targeted agents: Should we pay more attention to late toxicities? *Journal of Clinical Oncology*, 29(13):1728–1735.

- Prasad, V. and Mailankody, S. (2017). Research and development spending to bring a single cancer drug to market and revenues after approval. *JAMA Internal Medicine*, 177(11):1569–1575.
- Pratt, V., McLeod, H., Dean, L., Malheiro, A., and Rubinstein, W. (2015). *Medical Genetics Summaries*. NCBI Bookshelf.
- Proust-Lima, C., Séne, M., Taylor, J. M. G., and Jacqmin-Gadda, H. (2014). Joint latent class models for longitudinal and time-to-event data: A review. *Statistical Methods in Medical Research*, 23(1):74–90.
- Reitsma, D., Combest, A., Hummel, J., and Simmons, A. (2015). Improving Oncology Trials Through Adaptive Designs. *Applied Clinical Trials*.
- Riviere, M.-K., Yuan, Y., Jourdan, J.-H., Dubois, F., and Zohar, S. (2016). Phase I/II dose-finding design for molecularly targeted agent: Plateau determination using adaptive randomization. *Statistical Methods in Medical Research*.
- Rizopoulos, D. (2012). *Joint models for longitudinal and time-to-event data with applications in R*. Biostatistics Series, Chapman & Hall/CRC.
- Robertson, A., Wright, F., and Dykstra, R. (1988). *Order restricted statistical inference*. New York, Wiley.
- Rogatko, A., Babb, J. S., Tighiouart, M., Khuri, F. R., and Hudes, G. (2005). New paradigm in dose-finding trials: patient-specific dosing and beyond phase i. *Clin Cancer Res.*, 11(15):5342–5346.
- Rouanet, A., Helmer, C., Dartigues, J.-F., and Jacqmin-Gadda, H. (2017). Interpretation of mixed models and marginal models with cohort attrition due to death and drop-out. *Statistical Methods in Medical Research*, page 096228021772367.
- Rustin, G. J. S., Quinn, M., Thigpen, T., du Bois, A., Pujade-Lauraine, E., Jakobsen, A., Eisenhauer, E., Sagae, S., Greven, K., Vergote, I., Cervantes, A., and Vermorken, J. (2004). New Guidelines to Evaluate the Response to Treatment in Solid Tumors (Ovarian Cancer). *Journal of the National Cancer Institute*, 96(6):487–488.

- Rustin, G. J. S., Vergote, I., Eisenhauer, E., Pujade-Lauraine, E., Quinn, M., Thigpen, T., du Bois, A., Kristensen, G., Jakobsen, A., Sagae, S., Greven, K., Parmar, M., Friedlander, M., Cervantes, A., and Vermorken, J. (2011). Definitions for response and progression in ovarian cancer clinical trials incorporating RECIST 1.1 and CA 125 agreed by the Gynecological Cancer Intergroup (GCIG). *International journal of gynecological cancer : official journal of the International Gynecological Cancer Society*, 21(2):419–423.
- Schluchter, M. D. (1992). Methods for the analysis of informatively censored longitudinal data. *Statistics in Medicine*, 11(14):1861–1870.
- Sharma, P., Kumar, L., Mohanty, S., and Kochupillai, V. (2010). Response to Imatinib mesylate in chronic myeloid leukemia patients with variant BCR-ABL fusion transcripts. *Annals of hematology*, 89(3):241–247.
- Shen, L. Z. and O’Quigley, J. (1996). Consistency of continual reassessment method under under model misspecification. *Biometrika*, 83(2):395–405.
- Simon, R., Freidlin, B., Rubinstein, L., Arbuch, S. G., Collins, J., and Christian, M. C. (1997). Accelerated titration designs for phase I clinical trials in oncology. *J Nat Can Inst*, 89(15):1138–1147.
- Sinclair, K. and Whitehead, A. (2014). A Bayesian approach to dose-finding studies for cancer therapies: Incorporating later cycles of therapy. *Statistics in Medicine*, 33(15):2665–2680.
- Skolnik, J. M., Barrett, J. S., Jayaraman, B., Patel, D., and Adamson, P. C. (2008). Shortening the timeline of pediatric phase I trials: the rolling six design. *J Clin Oncol*, 26(2):190–195.
- Soria, J. C. (2011). Phase 1 trials of molecular targeted therapies: Are we evaluating toxicities properly? *European Journal of Cancer*, 47(10):1443–1445.
- Thall, P. F. and Cook, J. D. (2004). Dose-finding based on efficacy–toxicity trade-offs. *Biometrics*, 60:684–693.

- Tourneau, C. L., JackLee, J., and Siu, L. L. (2009). Dose escalation methods in phase i cancer clinical trials. *J Natl Cancer Inst*, 101(10):708–720.
- Tsiatis, A. A. and Davidian, M. (2004). Joint modeling of Longitudinal and Time-to-Event Data: An Overview. *Statistica Sinica*, 14:809–834.
- Wages, N. A. and Tait, C. (2015). Seamless phase I/II adaptive design for oncology trials of molecularly targeted agents. *Journal of biopharmaceutical statistics*, 25(5):903–920.
- Wikipedia (2018). *HER2/neu*. <https://en.wikipedia.org/wiki/HER2/neu> [Accessed: 20/05/2018].
- Wong, C. H., Siah, K. W., and Lo, A. W. (2018). Estimation of clinical trial success rates and related parameters. *Biostatistics*, (June):1–14.
- Wulfsohn, M. S. and Tsiatis, A. A. (1997). A joint model for survival and longitudinal data measured with error. *Biometrics*, 53(1):330–339.
- Yeung, W. Y., Reigner, B., Beyer, U., Diack, C., Sabanés bové, D., Palermo, G., and Jaki, T. (2017). Bayesian adaptive dose-escalation designs for simultaneously estimating the optimal and maximum safe dose based on safety and efficacy. *Pharmaceutical Statistics*, (March):1–18.
- Yeung, W. Y., Whitehead, J., Reigner, B., Beyer, U., Diack, C., and Jaki, T. (2015). Bayesian adaptive dose-escalation procedures for binary and continuous responses utilizing a gain function. *Pharmaceutical Statistics*, 14(6):479–487.
- Yuan, Y. and Yin, G. (2009). Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes. *Journal of the Royal Statistical Society*, 58:719–736.
- Zang, Y., Jack Lee, J., and Yuan, Y. (2015). Adaptive designs for identifying optimal biological dose for molecularly targeted agents. *Clin Trials.*, 11(3):319–327.
- Zhang, W., Sargent, D. J., and Mandrekar, S. (2006). An adaptive dose-finding design incorporating both toxicity and efficacy. *Statistics in Medicine*, 25(14):2365–2383.

- Zhang, X., Zhang, Y., Ye, X., Guo, X., Zhang, T., and He, J. (2016). Overview of phase IV clinical trials for postmarket drug safety surveillance: A status report from the ClinicalTrials.gov registry. *BMJ Open*, 6(11):1–9.
- Zhou, C., Clamp, A., Backen, A., Berzuini, C., Renehan, A., Banks, R. E., Kaplan, R., Scherer, S. J., Kristensen, G. B., Pujade-Lauraine, E., Dive, C., and Jayson, G. C. (2016). Systematic analysis of circulating soluble angiogenesis-associated proteins in ICON7 identifies Tie2 as a biomarker of vascular progression on bevacizumab. *British Journal of Cancer*, 115(2):228–235.

Appendix A

Table A.1: Simulation results of 2000 replicates. Shown are the percentage of bias and the coverage of the joint model parameters over six different sample sizes: $N = 15$, $N = 20$, $N = 25$, $N = 30$, $N = 40$, and $N = 60$ and for unbalanced data for the linear mixed effects model.

Parameter	Bias						Coverage					
	N=15	N=20	N=25	N=30	N=40	N=60	N=15	N=20	N=25	N=30	N=40	N=60
Longitudinal												
β_0	-1.18	-0.36	-0.16	-0.10	-0.13	-0.02	0.92	0.93	0.95	0.95	0.95	0.94
β_1	-0.18	0.02	0.00	0.00	-0.01	-0.01	0.91	0.93	0.94	0.93	0.94	0.95
β_2	0.58	-0.13	0.20	0.02	0.02	-0.08	0.89	0.91	0.91	0.92	0.93	0.94
β_3	-14.92	4.00	5.40	-0.25	-0.13	2.30	0.92	0.90	0.91	0.93	0.93	0.93
σ_1	-5.32	-8.20	-8.30	-7.30	-5.90	-6.20	0.90	0.99	0.99	0.99	0.99	0.98
σ_2	12.00	-3.40	-2.50	-2.10	-1.40	-0.88	0.98	0.97	0.97	0.97	0.96	0.96
Survival												
a_0	1.28	0.23	-0.19	0.15	-0.12	0.09	0.98	0.97	0.95	0.95	0.95	0.95
a_1	-27.42	13.50	8.80	7.50	6.01	3.10	0.98	0.97	0.97	0.96	0.96	0.95
γ	368.71	40.62	-25.37	12.29	10.96	-2.70	0.99	0.99	0.99	0.99	0.99	0.97

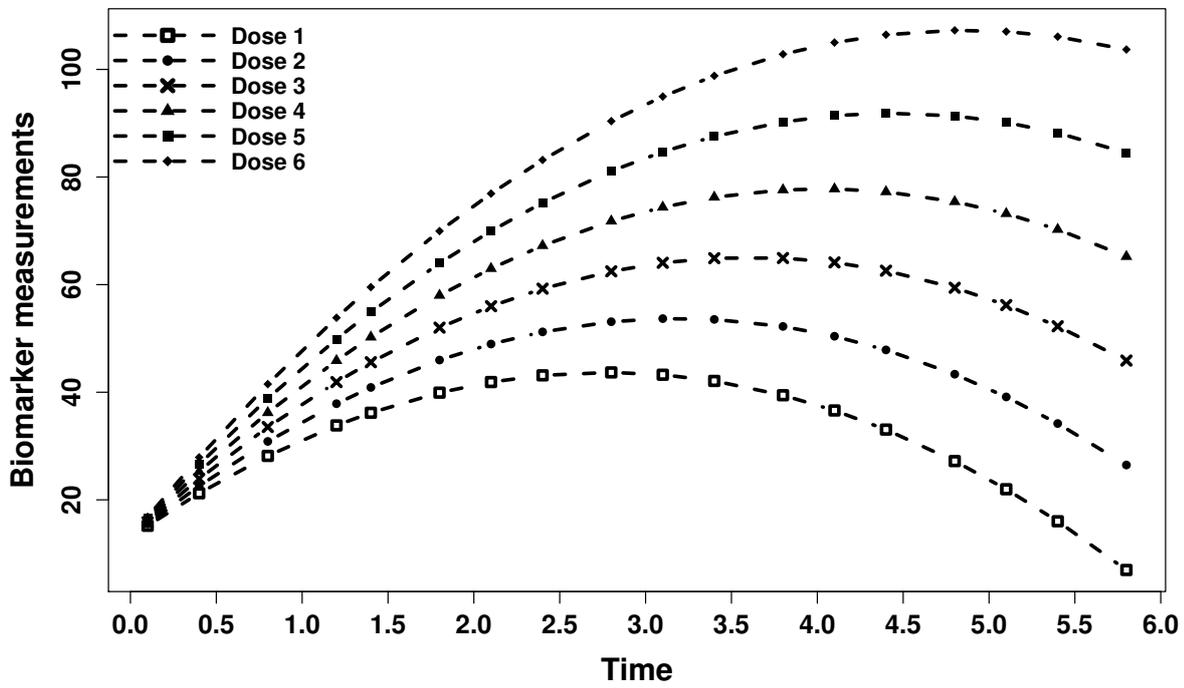


Figure A.1: Biomarker trajectory over time and for 6 dose levels.

Table A.2: Pocock-type sequential boundaries to monitor dose-limiting toxicity rate. Toxicity rate was set at 35% and the one-sided level test at 10%.

Number of patients, (N)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Boundary, (b_N)	-	-	3	4	4	5	5	5	6	6	7	7	8	8	8

Table A.3: Joint model parameters for data generation of the eight principal scenarios. Standard deviations were fixed at $\sigma_1 = 1$ and $\sigma_2 = 3$, and the parameter of the shared random effect at $\gamma = 0.1$.

Parameter	Scenario							
	1	2	3	4	5	6	7	8
Longitudinal								
β_0	270	90	115	115	80	270	115	90
β_1	35	8.5	15	15	7	35	10	8.5
β_2	-80	-19	-31.5	-31.5	-11	-80	-27	-19
β_3	-4	-4	-4	-4	-4	-4	-4	-4
Survival								
a_0	7.50	7.43	5.20	7.43	5.82	7.20	7.20	2.7
a_1	0.33	0.33	0.30	0.33	1.94	0.30	0.30	1.5
a_g	-3.30	-3.30	-3.50	-3.30	-3.12	-2.90	-2.90	-2.6

Table A.4: Sensitivity analysis of 2000 replicates and a sample size of 60. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}), under the scenarios of Table 2.3. Residual standard deviation was $\sigma_2 = 4$ and random effect standard deviation $\sigma_1 = 2$. The optimal dose is in bold and the MTD in italic.

Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
1	$(Y_{(t),min}, p_t)$	(200, 0.00)	(176, 0.00)	(147, 0.02)	(115, 0.12)	(79, 0.38)	(40, 0.75)	
	$P_{\%}$	0.0	0.0	0.0	1.0	97.8	1.2	0.0
	\bar{N}_{pat}	2.0	2.0	2.2	5.4	40.7	7.7	
2	$(Y_{(t),min}, p_t)$	(71, 0.00)	(64, 0.00)	(57, 0.03)	(49, 0.12)	<i>(39, 0.38)</i>	(29, 0.75)	
	$P_{\%}$	0.3	10.7	80.0	8.5	0.5	0.0	0.0
	\bar{N}_{pat}	2.3	7.2	31.4	8.2	5.8	5.1	
3	$(Y_{(t),min}, p_t)$	(87, 0.32)	(78, 0.71)	(67, 0.96)	(55, 0.99)	(41, 0.99)	(26, 0.99)	
	$P_{\%}$	99.4	0.2	0.0	0.0	0.0	0.0	0.4
	\bar{N}_{pat}	51.6	6.3	2.0	0.1	0.0	0.0	
4	$(Y_{(t),min}, p_t)$	(87, 0.00)	(78, 0.00)	(67, 0.03)	(55, 0.12)	<i>(41, 0.38)</i>	(26, 0.75)	
	$P_{\%}$	0.1	0.0	3.9	91.5	4.5	0.0	0.0
	\bar{N}_{pat}	2.0	2.1	5.5	36.3	9.0	5.0	
5	$(Y_{(t),min}, p_t)$	(71, 0.02)	(68, 0.07)	(65, 0.20)	<i>(61, 0.43)</i>	(57, 0.68)	(52, 0.88)	
	$P_{\%}$	99.6	0.2	0.0	0.2	0.0	0.0	0.0
	\bar{N}_{pat}	29.4	3.6	6.3	11.5	7.8	1.4	
6	$(Y_{(t),min}, p_t)$	(200, 0.00)	(176, 0.00)	(147, 0.01)	(115, 0.04)	(79, 0.16)	(40, 0.41)	
	$P_{\%}$	0.2	0.0	0.0	0.0	5.3	94.5	0.0
	\bar{N}_{pat}	2.0	2.0	2.0	2.4	8.1	43.4	
7	$(Y_{(t),min}, p_t)$	(84, 0.00)	(73, 0.00)	(61, 0.01)	(48, 0.04)	(32, 0.16)	<i>(16, 0.41)</i>	
	$P_{\%}$	0.0	0.0	0.0	14.9	84.9	0.2	0.0
	\bar{N}_{pat}	2.0	2.0	2.2	9.3	34.2	10.3	

Table A.5: Sensitivity analyses of 2000 replicates. Percentage of dose selection at the end of the trial ($P\%$) and mean number of patients assigned to each dose level (\bar{N}_{pat}), under the scenarios of Table 2.3, with different standard deviations, sample size, and random effects' distributions.

Conditions	Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
N=40									
$\sigma_1 = 2$ and $\sigma_2 = 4$	1	$P\%$	0.0	0.0	0.0	3.6	90.8	5.6	0.0
		\bar{N}_{pat}	2.0	2.0	2.2	5.1	21.7	7.0	
	2	$P\%$	0.4	11.9	74.4	11.9	1.4	0.0	0.0
		\bar{N}_{pat}	2.2	4.8	16.0	6.6	5.5	4.9	
	3	$P\%$	99.5	0.2	0.0	0.0	0.0	0.0	0.3
		\bar{N}_{pat}	31.9	6.0	1.9	0.2	0.0	0.0	
	4	$P\%$	0.0	0.0	6.9	83.1	10.0	0.0	0.0
		\bar{N}_{pat}	2.0	2.1	4.6	18.6	7.8	4.9	
	5	$P\%$	98.8	0.2	0.0	1.0	0.0	0.0	0.0
		\bar{N}_{pat}	16.7	3.5	5.3	8.2	5.3	1.0	
	6	$P\%$	0.2	0.0	0.1	0.0	11.5	88.2	0.0
		\bar{N}_{pat}	2.0	2.0	2.0	2.4	6.9	24.7	
	7	$P\%$	0.2	0.0	0.1	17.6	81.0	1.1	0.0
		\bar{N}_{pat}	2.0	2.0	2.2	6.4	18.1	9.2	
<hr/>									
$U \sim \Gamma(2, 2)$ for linear model									
$U \sim N(0, 1)$ for probit model	1	$P\%$	0.1	0.0	0.0	1.4	97.0	1.5	0.0
		\bar{N}_{pat}	2.0	2.0	2.2	5.8	40.4	7.6	
	2	$P\%$	0.4	11.7	73.7	14.0	0.2	0.0	0.0
		\bar{N}_{pat}	2.3	7.2	30.0	9.6	5.8	5.1	
	3	$P\%$	99.1	0.3	0.1	0.0	0.0	0.0	0.5
		\bar{N}_{pat}	51.7	6.3	1.9	0.1	0.0	0.0	
	4	$P\%$	0.0	0.0	4.2	87.4	8.4	0.0	0.0
		\bar{N}_{pat}	2.1	2.1	5.6	34.7	10.3	5.2	
	5	$P\%$	99.0	0.0	0.0	1.0	0.0	0.0	0.0
		\bar{N}_{pat}	29.5	3.5	6.4	11.6	7.7	1.3	
	6	$P\%$	0.2	0.0	0.0	0.0	4.5	95.3	0.0
		\bar{N}_{pat}	2.0	2.0	2.0	2.4	8.1	43.5	
	7	$P\%$	0.2	0.0	0.0	12.2	86.6	1.0	0.0
		\bar{N}_{pat}	2.1	2.0	2.2	8.3	34.2	11.2	

Table A.6: Sensitivity analysis of 2000 replicates and a sample size of 60. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). For scenarios A.6.1 – A.6.3 data was generated from a saturated linear mixed effects model for activity and the probit model for toxicity. For scenarios A.6.4 – A.6.6 data was generated from the linear mixed effects model for activity and a logistic model for toxicity. The optimal dose is in bold and the MTD in italic.

Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
A.6.1	$(Y_{(t),min}, p_t)$	(86, 0.00)	(73, 0.00)	(69, 0.01)	(66, 0.04)	(64, 0.13)	<i>(58, 0.35)</i>	
	$P_{\%}$	4.1	91.1	4.3	0.5	0.0	0.0	0.0
	\bar{N}_{pat}	3.9	30.0	7.6	3.3	3.9	11.3	
A.6.2	$(Y_{(t),min}, p_t)$	(89, 0.02)	(88, 0.07)	(66, 0.20)	<i>(59, 0.43)</i>	(48, 0.68)	(33, 0.88)	
	$P_{\%}$	17.6	2.9	76.9	2.5	0.1	0.0	0.0
	\bar{N}_{pat}	7.4	3.7	27.2	12.5	7.8	1.4	
A.6.3	$(Y_{(t),min}, p_t)$	(97, 0.00)	(94, 0.00)	(81, 0.01)	(76, 0.04)	(74, 0.13)	<i>(69, 0.35)</i>	
	$P_{\%}$	0.2	8.5	59.0	32.0	0.2	0.1	0.0
	\bar{N}_{pat}	2.1	5.8	19.5	17.1	3.9	11.6	
A.6.4	$(Y_{(t),min}, p_t)$	(200, 0.04)	(176, 0.08)	(147, 0.13)	(115, 0.23)	(79, 0.36)	(40, 0.54)	
	$P_{\%}$	0.4	0.0	0.2	10.8	66.4	22.2	0.0
	\bar{N}_{pat}	2.3	2.4	3.6	9.8	27.4	14.5	
A.6.5	$(Y_{(t),min}, p_t)$	(78, 0.10)	(71, 0.19)	<i>(62, 0.33)</i>	(53, 0.51)	(42, 0.72)	(30, 0.88)	
	$P_{\%}$	81.6	17.6	0.4	0.2	0.2	0.0	0.0
	\bar{N}_{pat}	30.8	14.8	6.0	5.2	2.6	0.6	
A.6.6	$(Y_{(t),min}, p_t)$	(193, 0.25)	(167, 0.40)	(135, 0.59)	(99, 0.77)	(60, 0.91)	(16, 0.97)	
	$P_{\%}$	13.5	71.5	14.9	0.1	0.0	0.0	0.0
	\bar{N}_{pat}	12.4	32.9	12.2	2.1	0.4	0.0	

Table A.7: Sensitivity analysis of 2000 replicates and a sample size of 60. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). For the activity data was generated from a linear mixed effects model, with response being associated with increase and progression with decrease of the biomarker, as shown in Figure A.1. The optimal dose is in bold and the MTD in italic.

Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
A.7.1	$(Y_{(t),min}, p_l)$	(39, 0.02)	(43, 0.07)	(46, 0.20)	<i>(49, 0.43)</i>	(52, 0.68)	(75, 0.88)	
	$P_{\%}$	99.8	0.0	0.0	0.2	0.0	0.0	0.0
	\bar{N}_{pat}	32.9	3.6	6.2	10.5	6.0	0.8	
A.7.2	$(Y_{(t),min}, p_l)$	(44, 0.00)	(53, 0.00)	(62, 0.03)	(72, 0.12)	<i>(81, 0.38)</i>	(90, 0.75)	
	$P_{\%}$	0.2	6.7	73.7	19.0	0.4	0.0	0.0
	\bar{N}_{pat}	2.1	5.1	28.9	12.8	6.6	4.5	

Table A.8: Sensitivity analysis of 2000 replicates and a sample size of 60. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). For the activity data was generated from the linear mixed effects model assuming both a random intercept and a random slope for time. Toxicity model was not modified. The optimal dose is in bold and the MTD in italic.

Scenario	Conditions		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
		$(Y_{(t),min}, p_l)$	(71, 0.00)	(64, 0.00)	(57, 0.03)	(49, 0.12)	<i>(39, 0.38)</i>	(29, 0.75)	
A.8.1	$\sigma_1 = 1, \sigma_2 = 3$	$P_{\%}$	0.2	10.8	80.3	8.0	0.7	0.0	0.0
	$\sigma_3 = 1, \rho_{\sigma_1, \sigma_3} = 0.7$	\bar{N}_{pat}	2.2	7.4	31.1	8.4	5.8	5.0	
A.8.2	$\sigma_1 = 1, \sigma_2 = 3$	$P_{\%}$	1.0	11.8	67.9	18.2	1.2	0.0	0.0
	$\sigma_3 = 3, \rho_{\sigma_1, \sigma_3} = 0.7$	\bar{N}_{pat}	2.6	7.9	26.5	11.6	6.2	5.2	

Table A.9: Sensitivity analysis of 2000 replicates and a sample size of 60. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). For the activity data was generated from a linear mixed effects model that included both a linear and a quadratic term for time. Toxicity model was not modified. The optimal dose is in bold and the MTD in italic.

Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
A.9.1	$(Y_{(l),min}, p_l)$	(261, 0.00)	(223, 0.00)	(183, 0.03)	(138, 0.12)	(89, 0.38)	(37, 0.75)	
	$P_{\%}$	0.2	0.0	0.0	1.1	97.5	1.2	0.0
	\bar{N}_{pat}	2.1	2.0	2.2	5.2	40.6	7.9	

Appendix B

Table B.1: Joint model parameters for data generation of the five principal scenarios. Standard deviations were fixed at $\sigma_1 = 1$ and $\sigma_2 = 3$, and the parameter of the shared random effect at $\gamma = 0.1$.

Parameter	Scenario				
	1	2	3	4	5
Longitudinal					
β_0	160	200	90	115	90
β_1	8.5	9.5	12	16.5	12
β_2	-40	-40	-20	-30	-20
Survival					
a_0	5.82	7.44	5.11	5.82	2.70
a_1	1.94	0.33	1.94	1.94	1.50
a_g	-3.12	-3.30	-3.00	-3.12	-2.60

Table B.2: Sensitivity analyses of 1000 replicates. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}), under the scenarios of Table 3.1, with different standard deviations, sample size, and random effects' distributions.

Conditions	Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
$\sigma_1 = 2$ and $\sigma_2 = 5$									
N=40	1	$P_{\%}$	0.1	2.3	97.5	0.1	0.0	0.0	0.0
		\bar{N}_{pat}	2.1	3.3	20.6	7.9	5.3	1.0	
	2	$P_{\%}$	0.0	0.0	0.0	4.2	95.8	0.0	0.0
		\bar{N}_{pat}	2.0	2.0	2.2	5.3	23.6	4.9	
	3	$P_{\%}$	98.6	0.8	0.1	0.1	0.0	0.0	0.4
		\bar{N}_{pat}	18.6	5.2	8.3	6.2	1.6	0.1	
	4	$P_{\%}$	17.3	79.0	3.6	0.1	0.0	0.0	0.0
		\bar{N}_{pat}	5.1	15.0	5.8	8.0	5.2	0.9	
<hr/>									
$U \sim \Gamma(2, 2)$ for linear model									
$U \sim N(0, 1)$ for probit model									
N=40 and $\sigma_2 = 3$	1	$P_{\%}$	0.0	0.1	99.9	0.0	0.0	0.0	0.0
		\bar{N}_{pat}	2.1	2.5	21.1	8.2	5.1	1.0	
	2	$P_{\%}$	0.0	0.0	0.0	5.8	94.2	0.0	0.0
		\bar{N}_{pat}	2.0	2.0	2.2	5.3	23.5	5.0	
	3	$P_{\%}$	97.0	1.8	1.1	0.0	0.0	0.0	0.1
		\bar{N}_{pat}	18.7	5.3	8.2	6.2	1.5	0.1	
	4	$P_{\%}$	16.8	82.5	0.6	0.1	0.0	0.0	0.0
		\bar{N}_{pat}	5.1	15.5	5.4	8.0	5.1	1.0	

Table B.3: Sensitivity analyses of 1000 replicates and a sample size of 60. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). Toxicity data was generated assuming increasing hazard at each successive treatment cycle. The optimal dose is in bold, the MTD in italic and the beginning of the plateau is underlined.

Scenario		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
B.3.1	$(Y_{(l),min}, p_l)$	(88, 0.01)	<u>(59, 0.03)</u>	(59, 0.11)	<i>(59, 0.31)</i>	(59, 0.61)	(59, 0.87)	
	$P_{\%}$	0.0	100.0	0.0	0.0	0.0	0.0	0.0
	\bar{N}_{pat}	2.2	34.8	3.0	3.0	3.0	14.0	
B.3.2	$(Y_{(l),min}, p_l)$	(201, 0.00)	(176, 0.00)	(147, 0.00)	(114, 0.03)	(77, 0.14)	(38, 0.42)	
	$P_{\%}$	0.0	0.0	0.0	0.0	6.5	93.5	0.0
	\bar{N}_{pat}	2.0	2.0	2.1	2.4	7.9	43.6	

Table B.4: Sensitivity analyses of 1000 replicates and a sample size of 60. Percentage of dose selection at the end of the trial ($P_{\%}$) and mean number of patients assigned to each dose level (\bar{N}_{pat}). For the activity data was generated from the log model 2.2, assuming both a random intercept and a random slope for time. Toxicity model was not modified. The optimal dose is in bold and the MTD in italic.

Model		Dose 1	Dose 2	Dose 3	Dose 4	Dose 5	Dose 6	None Selected
	$(Y_{(l),min}, p_l)$	(71, 0.00)	(64, 0.00)	(57, 0.03)	(49, 0.12)	<i>(39, 0.38)</i>	(29, 0.75)	
$\sigma_1 = 1, \sigma_2 = 3$								
$\sigma_3 = 1, \rho_{\sigma_1, \sigma_3} = 0.7$								
log	$P_{\%}$	0.2	10.8	80.3	8.0	0.7	0.0	0.0
	\bar{N}_{pat}	2.2	7.4	31.1	8.4	5.8	5.0	
plateau	$P_{\%}$	0.6	14.2	78.9	5.3	0.1	0.0	0.0
	\bar{N}_{pat}	2.4	8.4	31.2	6.3	6.5	5.2	
$\sigma_1 = 1, \sigma_2 = 3$								
$\sigma_3 = 3, \rho_{\sigma_1, \sigma_3} = 0.7$								
log	$P_{\%}$	1.0	11.8	67.9	18.2	1.2	0.0	0.0
	\bar{N}_{pat}	2.6	7.9	26.5	11.6	6.2	5.2	
plateau	$P_{\%}$	0.7	15.9	73.9	7.3	2.2	0.0	0.0
	\bar{N}_{pat}	2.4	9.1	29.3	7.7	6.6	4.8	

Appendix C

C.1 Descriptives

Table C.1: Number (#) and percentage (%) of patients per treatment cycle and dose subgroup.

	cycle 1	cycle 2	cycle 3	cycle 4	cycle 5	cycle 6
# patients d<MTD	490	338	157	104	60	43
% patients d<MTD	100	69.0	32.0	21.2	12.2	8.8
# patients d=MTD	289	152	54	37	25	20
% patients d=MTD	100	52.6	18.7	12.8	8.7	6.9
# patients d>MTD	163	58	25	14	8	5
% patients d>MTD	100	35.6	15.3	8.6	4.9	3.1

Table C.2: Number (#) and percentage (%) of patients per type of toxicity and treatment cycle.

	cycle 1	cycle 2	cycle 3	cycle 4	cycle 5	cycle 6
# non-hematologic	142	33	19	4	4	1
% non-hematologic	15.1	6.0	8.1	2.6	4.3	1.5
# hematologic	63	18	2	6	1	1
% hematologic	6.7	3.3	0.9	3.9	1.1	1.5
# both*	33	2	2	0	1	1
% both	3.5	0.4	0.9	0.000	1.1	1.5

* Both non-hematologic and hematologic toxicity at the same cycle

C.2 Main analysis

Table C.3: Association of severe toxicity with cycle, for groups of patients treated at doses below, above and at the MTD.

Group of dose (n)*	Parameter	Estimate	S.E.**	Lower CI***	Upper CI
below the MTD (490)	intercept	1.171	0.065	1.045	1.298
	cycle	0.160	0.044	0.074	0.246
at the MTD (289)	intercept	0.603	0.072	0.461	0.745
	cycle	0.307	0.062	0.187	0.428
above the MTD (163)	intercept	0.027	0.092	-0.152	0.207
	cycle	0.359	0.079	0.204	0.513

* Sample size

** Standard error

*** Confidence interval

Table C.4: Risk and cumulative incidence function (CIF) of severe toxicity over 6 cycles, for groups of patients treated at doses below, above and at the MTD.

Group of dose (n)*	Treatment cycle	Risk	Lower PI**	Upper PI	CIF	Lower PI	Upper PI
below the MTD (490)	1	0.121	0.097	0.152	0.121	0.097	0.152
	2	0.091	0.076	0.109	0.201	0.166	0.241
	3	0.068	0.049	0.087	0.255	0.216	0.300
	4	0.049	0.028	0.072	0.292	0.246	0.344
	5	0.035	0.016	0.062	0.317	0.263	0.377
	6	0.024	0.008	0.053	0.333	0.270	0.405
at the MTD (289)	1	0.273	0.226	0.321	0.273	0.226	0.321
	2	0.181	0.147	0.216	0.405	0.350	0.461
	3	0.112	0.068	0.155	0.472	0.411	0.537
	4	0.064	0.026	0.112	0.505	0.427	0.572
	5	0.033	0.008	0.080	0.522	0.434	0.597
	6	0.016	0.002	0.055	0.529	0.437	0.615
above the MTD (163)	1	0.489	0.404	0.567	0.489	0.404	0.567
	2	0.350	0.285	0.417	0.668	0.594	0.740
	3	0.228	0.127	0.316	0.744	0.667	0.813
	4	0.135	0.041	0.241	0.778	0.691	0.852
	5	0.072	0.010	0.175	0.794	0.700	0.874
	6	0.034	0.001	0.118	0.801	0.701	0.887

* Sample size

** Prediction interval

C.3 Complete model

Table C.5: Association of severe toxicity with cycle and dose.

Parameter	Estimate	S.E.*	Lower CI**	Upper CI
intercept	1.381	0.075	1.233	1.529
cycle	0.245	0.033	0.181	0.309
dose	-0.727	0.070	-0.864	-0.589

* Standard error

** Confidence interval

Table C.6: Risk and cumulative incidence function (CIF) of severe toxicity over 6 cycles, for patients treated at the MTD.

Treatment	Risk	Lower PI*	Upper PI	CIF	Lower PI	Upper PI
cycle						
1	0.256	0.227	0.286	0.256	0.226	0.285
2	0.184	0.164	0.206	0.394	0.359	0.431
3	0.126	0.103	0.152	0.470	0.433	0.512
4	0.083	0.058	0.111	0.514	0.470	0.561
5	0.051	0.029	0.081	0.539	0.489	0.593
6	0.030	0.013	0.056	0.553	0.499	0.614

* Prediction interval

C.4 Residual Check

C.4.1 For cycle effect

Table C.7: Association of severe toxicity with cycle. For the cycle variable, we assume in model 1 a different effect among cycles 1-3 and 4-6 and in model 2 a different effect among cycles 1-2 and 3-6.

	Parameter	Estimate	S.E.*	Lower CI**	Upper CI
Model 1	intercept	0.570	0.076	0.422	0.719
	cycle 1-3	0.409	0.101	0.210	0.608
	cycle 4-6	0.270	0.065	0.143	0.396
Model 2	intercept	0.525	0.077	0.375	0.676
	cycle 1-2	0.691	0.155	0.388	0.995
	cycle 3-6	0.276	0.058	0.162	0.390

* Standard error

** Confidence interval

Table C.8: Risk and cumulative incidence function of severe toxicity assuming a) 1 linear cycle effect b) 2 cycle effects 1-3 and 4-6 and c) 2 cycle effects 1-2 and 3-6.

Treatment cycle	Cycles 1-6 linear	Cycles 1-3 Cycles 4-6	Lower PI*	Upper PI	Cycles 1-2 Cycles 3-6	Lower PI	Upper PI
Risk of severe toxicity							
1	0.273	0.284	0.228	0.334	0.300	0.247	0.353
2	0.181	0.164	0.119	0.214	0.112	0.064	0.164
3	0.112	0.133	0.120	0.216	0.141	0.091	0.192
4	0.064	0.084	0.024	0.151	0.088	0.042	0.150
5	0.033	0.049	0.008	0.120	0.051	0.016	0.112
6	0.016	0.027	0.002	0.093	0.028	0.005	0.085
Cumulative incidence function							
1	0.273	0.284	0.228	0.334	0.300	0.247	0.353
2	0.405	0.401	0.343	0.455	0.378	0.320	0.433
3	0.472	0.451	0.377	0.525	0.465	0.406	0.534
4	0.505	0.497	0.417	0.574	0.512	0.441	0.593
5	0.522	0.522	0.433	0.608	0.537	0.454	0.632
6	0.529	0.535	0.436	0.636	0.551	0.460	0.656

* Prediction interval

C.4.2 For dose effect

Table C.9: Association of severe toxicity with cycle and dose. For the dose we assume 2 dose effects; 1 for patients treated below the MTD and 1 for patients treated above or at the MTD.

Parameter	Estimate	S.E.*	Lower CI**	Upper CI
intercept	1.103	0.058	0.990	1.217
cycle	0.241	0.033	0.177	0.306
dose1	-0.457	0.080	-0.614	-0.300
dose2	-1.004	0.094	-1.188	-0.820

* Standard error

** Confidence interval

Table C.10: Risk and cumulative incidence function of severe toxicity assuming 2 variables for the dose effect.

Treatment cycle	Dose linear	Dose 1 - Dose 2	Lower PI*	Upper PI
Risk of severe toxicity				
1	0.256	0.259	0.216	0.301
2	0.184	0.187	0.155	0.220
3	0.126	0.129	0.097	0.161
4	0.083	0.085	0.055	0.118
5	0.051	0.053	0.027	0.084
6	0.030	0.032	0.013	0.060
Cumulative incidence function				
1	0.256	0.259	0.216	0.301
2	0.394	0.398	0.340	0.453
3	0.470	0.476	0.410	0.536
4	0.514	0.521	0.450	0.586
5	0.539	0.546	0.471	0.617
6	0.553	0.561	0.480	0.638

* Prediction interval

C.5 Competing events

C.5.1 Naive analysis

Table C.11: Association of severe toxicity with cycle for patients treated at doses below, above and at the MTD. Estimation of hazards per type of toxicity.

Group of dose (n)*	Parameter	Estimate	S.E.**	Lower CI***	Upper CI
non-hematologic toxicity					
below the MTD (490)	intercept	1.383	0.072	1.243	1.523
	cycle	0.151	0.048	0.056	0.246
at the MTD (289)	intercept	0.909	0.076	0.759	1.058
	cycle	0.162	0.051	0.061	0.263
above the MTD (163)	intercept	0.282	0.092	0.103	0.462
	cycle	0.297	0.071	0.159	0.436
hematologic toxicity					
below the MTD (490)	intercept	1.678	0.085	1.511	1.846
	cycle	0.121	0.057	0.010	0.232
at the MTD (289)	intercept	1.214	0.089	1.040	1.388
	cycle	0.256	0.076	0.106	0.405
above the MTD (163)	intercept	0.895	0.104	0.690	1.099
	cycle	0.264	0.082	0.103	0.425

* Sample size

** Standard error

*** Confidence interval

Table C.12: Cumulative incidence function (CIF) of severe toxicity for patients treated at doses below, above and at the MTD. CIFs were calculated for non-hematologic toxicity alone and and hematologic toxicity alone.

Group of dose (n)*	Treatment cycle	CIF	Lower PI**	Upper PI	CIF	Lower PI	Upper PI
		non-hematologic			hematologic		
below the MTD (490)	1	0.083	0.062	0.107	0.047	0.032	0.066
	2	0.141	0.112	0.173	0.081	0.059	0.109
	3	0.180	0.146	0.219	0.106	0.078	0.140
	4	0.207	0.165	0.252	0.125	0.090	0.166
	5	0.226	0.176	0.281	0.138	0.094	0.186
	6	0.238	0.182	0.307	0.148	0.099	0.207
at the MTD (289)	1	0.182	0.141	0.225	0.112	0.082	0.150
	2	0.298	0.246	0.356	0.175	0.133	0.223
	3	0.374	0.314	0.438	0.210	0.164	0.267
	4	0.425	0.354	0.499	0.229	0.173	0.299
	5	0.460	0.376	0.548	0.239	0.177	0.318
	6	0.483	0.389	0.590	0.243	0.178	0.331
above the MTD (163)	1	0.389	0.320	0.472	0.185	0.128	0.248
	2	0.561	0.488	0.641	0.286	0.216	0.357
	3	0.644	0.568	0.723	0.341	0.263	0.440
	4	0.687	0.598	0.770	0.371	0.276	0.475
	5	0.709	0.609	0.801	0.388	0.275	0.510
	6	0.720	0.612	0.822	0.396	0.277	0.536

* Sample size

** Prediction interval

C.5.2 Competing risks analysis

Table C.13: Association of severe toxicity with cycle for patients treated at doses below, above and at the MTD. Estimation of cause-specific hazards per type of toxicity.

Group of dose (n)*	Parameter	Estimate	S.E.**	Lower CI***	Upper CI
non-hematologic toxicity					
below the MTD (490)	intercept	1.437	0.074	1.291	1.583
	cycle	0.153	0.052	0.051	0.255
at the MTD (289)	intercept	0.985	0.081	0.827	1.143
	cycle	0.211	0.065	0.084	0.338
above the MTD (163)	intercept	0.515	0.098	0.323	0.707
	cycle	0.367	0.103	0.165	0.570
both hematologic and non-hematologic toxicity					
below the MTD (490)	intercept	1.688	0.086	1.518	1.857
	cycle	0.121	0.059	0.006	0.236
at the MTD (289)	intercept	1.206	0.092	1.026	1.385
	cycle	0.356	0.103	0.155	0.557
above the MTD (163)	intercept	0.876	0.105	0.670	1.082
	cycle	0.177	0.088	0.004	0.351

* Sample size

** Standard error

*** Confidence interval

Table C.14: Cumulative incidence function (CIF) of severe toxicity for patients treated at doses below, above and at the MTD. CIFs were calculated for overall toxicity, non-hematologic toxicity and both hematologic and non-hematologic toxicity.

Group of dose(n)*	Treatment cycle	CIF overall	CIF non-hematologic	Lower PI**	Upper PI	CIF both	Lower PI	Upper PI
below the MTD (490)	1	0.121	0.075	0.056	0.099	0.046	0.030	0.064
	2	0.201	0.124	0.096	0.156	0.077	0.053	0.100
	3	0.255	0.157	0.126	0.193	0.098	0.070	0.126
	4	0.292	0.178	0.143	0.221	0.113	0.080	0.147
	5	0.317	0.193	0.151	0.243	0.124	0.086	0.167
	6	0.333	0.202	0.152	0.261	0.131	0.088	0.187
at the MTD (289)	1	0.273	0.162	0.121	0.201	0.114	0.080	0.150
	2	0.405	0.246	0.196	0.299	0.157	0.119	0.204
	3	0.472	0.294	0.235	0.354	0.173	0.128	0.224
	4	0.505	0.322	0.253	0.394	0.179	0.132	0.234
	5	0.522	0.338	0.261	0.419	0.182	0.133	0.241
	6	0.529	0.348	0.266	0.441	0.182	0.128	0.239
above the MTD (163)	1	0.489	0.303	0.238	0.375	0.190	0.129	0.251
	2	0.668	0.400	0.329	0.477	0.265	0.200	0.337
	3	0.744	0.435	0.356	0.515	0.301	0.229	0.378
	4	0.778	0.449	0.366	0.532	0.322	0.241	0.405
	5	0.794	0.454	0.362	0.543	0.334	0.246	0.426
	6	0.801	0.456	0.364	0.551	0.342	0.248	0.445

* Sample size

** Prediction interval

C.6 Sensitivity analysis

C.6.1 Reclassification of mixed cases with non-hematologic toxicities

Table C.15: Association of severe toxicity with cycle for patients treated at doses below, above and at the MTD. Estimation of cause-specific hazards per type of toxicity.

Group of dose (n)*	Parameter	Estimate	S.E.**	Lower CI***	Upper CI
both hematologic and non-hematologic toxicity					
below the MTD (490)	intercept	1.385	0.072	1.243	1.527
	cycle	0.164	0.051	0.063	0.264
at the MTD (289)	intercept	0.876	0.078	0.724	1.029
	cycle	0.241	0.064	0.115	0.367
above the MTD (163)	intercept	0.268	0.093	0.086	0.450
	cycle	0.340	0.083	0.177	0.503
hematologic toxicity					
below the MTD (490)	intercept	1.773	0.091	1.594	1.952
	cycle	0.103	0.060	-0.015	0.221
at the MTD (289)	intercept	1.373	0.100	1.178	1.568
	cycle	0.325	0.110	0.109	0.542
above the MTD (163)	intercept	1.306	0.128	1.055	1.557
	cycle	0.153	0.118	-0.079	0.384

* Sample size

** Standard error

*** Confidence interval

Table C.16: Cumulative incidence function (CIF) of severe toxicity for patients treated at doses below, above and at the MTD. CIFs were calculated for overall toxicity, both hematologic and non-hematologic toxicity and hematologic toxicity alone.

Group of dose (n)*	Treatment cycle	CIF	CIF both	Lower PI**	Upper PI	CIF	Lower PI	Upper PI
		overall				hematologic		
below the MTD (490)	1	0.121	0.083	0.069	0.119	0.038	0.024	0.057
	2	0.201	0.136	0.119	0.187	0.065	0.044	0.089
	3	0.255	0.171	0.151	0.231	0.084	0.059	0.111
	4	0.292	0.194	0.169	0.259	0.098	0.067	0.130
	5	0.317	0.208	0.178	0.277	0.108	0.074	0.148
	6	0.333	0.218	0.181	0.290	0.116	0.075	0.165
at the MTD (289)	1	0.273	0.190	0.147	0.234	0.085	0.053	0.119
	2	0.405	0.286	0.235	0.341	0.117	0.083	0.158
	3	0.472	0.338	0.278	0.400	0.130	0.091	0.181
	4	0.505	0.367	0.300	0.439	0.135	0.094	0.192
	5	0.522	0.383	0.308	0.466	0.137	0.094	0.192
	6	0.529	0.392	0.312	0.489	0.138	0.092	0.196
above the MTD (163)	1	0.489	0.394	0.316	0.480	0.096	0.057	0.144
	2	0.668	0.533	0.455	0.614	0.133	0.083	0.191
	3	0.744	0.590	0.508	0.681	0.150	0.098	0.215
	4	0.778	0.615	0.524	0.715	0.160	0.100	0.231
	5	0.794	0.627	0.530	0.721	0.167	0.102	0.243
	6	0.801	0.632	0.524	0.732	0.171	0.103	0.255

* Sample size

** Prediction interval

C.6.2 Exclusion of grade 3 hematologic toxicity

Table C.17: Number (#) and percentage (%) of patients per type of toxicity and treatment cycle.

	cycle 1	cycle 2	cycle 3	cycle 4	cycle 5	cycle 6
# non-hematologic	163	41	22	8	5	3
% non-hematologic	17.3	7.0	8.4	4.6	4.7	3.8
# hematologic	9	4	3	1	1	0
% hematologic	1.0	0.7	1.2	0.6	0.9	0.0
# both*	12	0	2	1	1	1
% both	1.3	0.0	0.8	0.6	0.9	1.3

* Both non-hematologic and hematologic toxicity at the same cycle

Table C.18: Association of severe toxicity with cycle, for groups of patients treated at doses below, above and at the MTD.

Group of dose (n)*	Parameter	Estimate	S.E.**	Lower CI***	Upper CI
below the MTD (490)	intercept	1.348	0.070	1.210	1.485
	cycle	0.144	0.047	0.052	0.235
at the MTD (289)	intercept	0.850	0.075	0.703	0.996
	cycle	0.157	0.050	0.058	0.256
above the MTD (163)	intercept	0.257	0.091	0.078	0.436
	cycle	0.273	0.071	0.135	0.411

* Sample size

** Standard error

*** Confidence interval

Table C.19: Risk and cumulative incidence function (CIF) of severe toxicity over 6 cycles, for groups of patients treated at doses below, above and at the MTD.

Group of dose (n)*	Treatment cycle	Risk	Lower CI**	Upper CI	CIF	Lower CI	Upper CI
below the MTD (490)	1	0.089	0.068	0.113	0.089	0.068	0.113
	2	0.068	0.055	0.083	0.151	0.121	0.183
	3	0.051	0.035	0.067	0.194	0.159	0.233
	4	0.038	0.021	0.057	0.225	0.182	0.270
	5	0.027	0.011	0.050	0.246	0.194	0.302
	6	0.019	0.005	0.044	0.260	0.201	0.329
at the MTD (289)	1	0.198	0.157	0.241	0.198	0.157	0.241
	2	0.157	0.129	0.188	0.324	0.269	0.376
	3	0.122	0.084	0.157	0.406	0.343	0.464
	4	0.093	0.048	0.137	0.462	0.385	0.530
	5	0.070	0.025	0.121	0.499	0.409	0.581
	6	0.051	0.012	0.108	0.525	0.420	0.622
above the MTD (163)	1	0.399	0.323	0.475	0.399	0.323	0.475
	2	0.298	0.245	0.362	0.578	0.504	0.662
	3	0.211	0.136	0.278	0.667	0.591	0.739
	4	0.141	0.055	0.222	0.714	0.625	0.792
	5	0.089	0.020	0.179	0.739	0.640	0.825
	6	0.052	0.005	0.143	0.753	0.646	0.845

* Sample size

** Prediction interval

Table C.20: Association of severe toxicity with cycle for patients treated at doses below, above and at the MTD. Estimation of cause-specific hazards per type of toxicity.

Group of dose (n)*	Parameter	Estimate	S.E.**	Lower CI***	Upper CI
non-hematologic toxicity					
below the MTD (490)	intercept	1.386	0.072	1.245	1.527
	cycle	0.160	0.050	0.063	0.258
at the MTD (289)	intercept	0.939	0.077	0.788	1.090
	cycle	0.161	0.053	0.057	0.265
above the MTD (163)	intercept	0.371	0.094	0.187	0.554
	cycle	0.327	0.082	0.168	0.487
both hematologic and non-hematologic toxicity					
below the MTD (490)	intercept	2.489	0.169	2.158	2.819
	cycle	-0.003	0.090	-0.180	0.174
at the MTD (289)	intercept	1.979	0.142	1.701	2.256
	cycle	0.058	0.092	-0.123	0.239
above the MTD (163)	intercept	1.667	0.150	1.374	1.961
	cycle	0.005	0.096	-0.182	0.192

* Sample size

** Standard error

*** Confidence interval

Table C.21: Cumulative incidence function (CIF) for patients treated at doses below, above and at the MTD. CIFs were calculated for overall toxicity, non-hematologic toxicity and both hematologic and non-hematologic toxicity.

Group of dose (n)*	Treatment cycle	CIF overall	CIF non-hematologic	Lower PI**	Upper PI	CIF both	Lower PI	Upper PI
below the MTD (490)	1	0.089	0.083	0.060	0.105	0.006	0.001	0.014
	2	0.151	0.138	0.109	0.173	0.012	0.003	0.023
	3	0.194	0.176	0.141	0.215	0.018	0.005	0.033
	4	0.225	0.201	0.158	0.247	0.023	0.009	0.045
	5	0.246	0.217	0.168	0.273	0.028	0.011	0.057
	6	0.260	0.228	0.173	0.295	0.033	0.012	0.075
at the MTD (289)	1	0.198	0.174	0.133	0.214	0.024	0.009	0.042
	2	0.324	0.283	0.232	0.341	0.041	0.019	0.067
	3	0.406	0.353	0.293	0.419	0.053	0.027	0.084
	4	0.462	0.399	0.328	0.472	0.062	0.031	0.100
	5	0.499	0.429	0.346	0.517	0.070	0.034	0.118
	6	0.525	0.450	0.356	0.550	0.075	0.032	0.140
above the MTD (163)	1	0.399	0.355	0.287	0.434	0.048	0.018	0.084
	2	0.578	0.501	0.428	0.578	0.076	0.037	0.123
	3	0.667	0.566	0.487	0.655	0.096	0.052	0.148
	4	0.714	0.595	0.509	0.692	0.111	0.059	0.180
	5	0.739	0.608	0.518	0.702	0.124	0.064	0.196
	6	0.753	0.614	0.521	0.715	0.136	0.064	0.228

* Sample size

** Prediction interval

Titre : Méthodes statistiques pour les essais de phase I/II de thérapies moléculaires ciblées en cancérologie.

Mots clés : Dose optimale ; Mesures de biomarqueur ; Modèles conjoints ; Recherche de dose ; Thérapies moléculaires ciblées ; Toxicité cumulative.

Résumé : Les essais cliniques de phase I en cancérologie permettent d'identifier la dose optimale (DO), définie comme la dose maximale tolérée (DMT). Les approches conventionnelles de recherche de dose reposent uniquement sur les événements de toxicité observés au cours du premier cycle de traitement. Le développement des thérapies moléculaires ciblées (TMC), habituellement administrées sur de longues périodes, a remis en question cet objectif. Considérer uniquement le premier cycle de traitement n'est pas suffisant. De plus, comme l'activité n'augmente pas nécessairement de façon monotone avec la dose, la toxicité et l'activité doivent être prises en compte pour identifier la DO. Récemment, les biomarqueurs continus sont de plus en plus utilisés pour mesurer l'activité. L'objectif de cette thèse était de proposer et d'évaluer des designs adaptatifs pour identifier la DO. Nous avons développé deux designs de recherche de dose, basés sur une modélisation conjointe des mesures longitudinales de l'activité des biomarqueurs et de la première toxicité dose-limitante (DLT), avec un effet aléatoire partagé. En utilisant des propriétés de distribution normales asymétriques,

l'estimation reposait sur la vraisemblance sans approximation ce qui est une propriété importante dans le cas de petits échantillons qui sont souvent disponibles dans ces essais. La DMT est associée à un certain risque cumulé de DLT sur un nombre prédéfini de cycles de traitement. La DO a été définie comme la dose la moins toxique parmi les doses actives, sous la contrainte de ne pas dépasser la DMT. Le second design étendait cette approche pour les cas d'une relation dose-activité qui pouvait atteindre un plateau. Un modèle à changement de pente a été implémenté. Nous avons évalué les performances des designs avec des études de simulations en étudiant plusieurs scénarios et divers degrés d'erreur de spécification des modèles. Finalement, nous avons effectué une analyse de 27 études des TMCs de phase I, en tant que monothérapie. Les études ont été réalisées par l'Institut National du Cancer. L'objectif principal était d'estimer le risque par cycle et l'incidence cumulative de la toxicité sévère, jusqu'à six cycles. Les analyses ont été effectuées séparément pour différents sous-groupes de doses, ainsi que pour les toxicités hématologiques et non-hématologiques.

Title : Statistical methods for phase I/II trials of molecularly targeted agents in oncology.

Keywords : Biomarker measurements ; Cumulative toxicity ; Dose-finding ; Joint modeling ; Molecularly targeted agents ; Optimal dose.

Abstract : Conventional dose-finding approaches in oncology of phase I clinical trials aim to identify the optimal dose (OD) defined as the maximum tolerated dose (MTD), based on the toxicity events observed during the first treatment cycle. The constant development of molecularly targeted agents (MTAs), usually administered in chronic schedules, has challenged this objective. Not only, the outcomes after the first cycle are of importance, but also activity does not necessarily increase monotonically with dose. Therefore, both toxicity and activity should be considered for the identification of the OD. Lately, continuous biomarkers are used more and more to monitor activity. The aim of this thesis was to propose and evaluate adaptive designs for the identification of the OD. We developed two dose-finding designs, based on a joint modeling of longitudinal continuous biomarker activity measurements and time to first dose limiting toxicity (DLT), with a shared random effect, using skewed normal distribution properties. Estimation relied on likelihood that did not require approximation, an important property in the context of small

sample sizes, typical of phase I/II trials. We addressed the important case of missing at random data that stem from unacceptable toxicity, lack of activity and rapid deterioration of phase I patients. The MTD was associated to some cumulative risk of DLT over a predefined number of treatment cycles. The OD was defined as the lowest dose within a range of active doses, under the constraint of not exceeding the MTD. The second design extended this approach for cases of a dose-activity relationship that could reach a plateau. A change point model was implemented. The performance of the approaches was evaluated through simulation studies, investigating a wide range of scenarios and various degrees of data misspecification. As a last part, we performed an analysis of 27 phase I studies of MTAs, as monotherapy, conducted by the National Cancer Institut. The primary focus was to estimate the per-cycle risk and the cumulative incidence function of severe toxicity, over up to six cycles. Analyses were performed separately for different dose subgroups, as well as for hematologic and non-hematologic toxicities.

