

**THESE DE DOCTORAT DE
L'UNIVERSITE PARIS-DIDEROT**

"ECOLE DOCTORALE PIERRE LOUIS DE SANTE PUBLIQUE A PARIS :
EPIDEMIOLOGIE ET SCIENCES DE L'INFORMATION BIOMEDICALE"

Spécialité: Biostatistiques

Présentée par
Marie-Karelle RIVIERE - JOURDAN

Pour obtenir le grade de:
Docteur de l'université Paris-Diderot

**DESIGNS ADAPTIFS DE RECHERCHE DE DOSE EN
ONCOLOGIE DANS LE CADRE DE COMBINAISONS
DE MOLECULES ET DE MOLECULES CIBLEES**

Directrice de Thèse: Sarah ZOHAR

Soutenue le: **12 Novembre 2014**

Devant le jury:

Présidente de jury :	France Mentré	Professeur
Rapporteurs :	Rodolphe Thiébaud	Professeur
	Jean-Michel Marin	Professeur
Examineurs :	François Coquet	Professeur
	Ying Yuan	Associate Professor
	William Rosenberger	Professeur
Directrice de thèse :	Sarah Zohar	Docteur
Responsable scientifique :	Frédéric Dubois	

Remerciements

Je tiens tout d'abord à remercier ma directrice de thèse, Sarah Zohar, pour son encadrement, ses encouragements et ses conseils qui m'ont fait avancer et progresser tout au long de ces trois années. Son enthousiasme m'a montré que la recherche est un univers passionnant. Je lui suis très reconnaissante pour le travail que nous avons accompli ensemble.

Je remercie Frédéric Dubois, mon responsable scientifique au sein d'IRIS, pour m'avoir encadrée et conseillée durant ces trois années. Son expérience sur les essais cliniques en cancérologie a permis de donner une dimension appliquée à mon travail et s'est avérée très enrichissante.

Mes remerciements vont également à Maylis Coste, Christine Gabarroca, ainsi que toute la « division méthodologie et valorisation des données » pour m'avoir accueillie dans leur entreprise, pour m'avoir permis d'avoir la chance de faire cette thèse et pour m'avoir impliquée dans leurs essais cliniques qui m'ont beaucoup appris.

Je remercie Anita Burgun et tous mes collègues de l'équipe 22 « Systèmes d'information et médecine personnalisée » de l'unité INSERM 1138 pour m'avoir accueillie dans leur équipe et pour toutes les conversations stimulantes que nous avons eues.

Je remercie le Professeur France Mentré qui m'a fait l'honneur de présider mon jury de thèse, le Professeur Marin et le Professeur Thiébaud qui ont accepté d'être les rapporteurs de ma thèse et ont contribué par leurs remarques et suggestions à améliorer la qualité de ce manuscrit. Je remercie également les Professeurs François Coquet, William Rosenberger et le Docteur Ying Yuan pour avoir accepté de participer à mon jury de thèse.

Je remercie l'ANRT et IRIS qui ont financé mon travail en m'accordant une bourse de thèse pour ces trois années.

Je remercie du fond du cœur ma famille, en particulier mes parents, mon frère, et ma tante, pour leur confiance, leur tendresse, et leur amour qui ont fait de moi ce que je suis aujourd'hui. Merci à ma Maman et à Catherine d'avoir relu la veille jusque très tard le résumé en français.

Je remercie tous mes amis pour leur soutien et pour tous les bons moments que nous avons passés ensemble.

Je tiens particulièrement à remercier mon mari pour son amour qui me porte et me guide tous les jours. Merci de m'entourer, merci d'avoir toujours cru en moi. Enfin merci pour sa patience à mes nombreuses questions sur le langage C, et surtout pour sa participation précieuse à la réalisation de mes packages R.

Contents

List of Tables	iii
List of Figures	v
Scientific production	vii
Glossary	ix
1 Introduction	3
1.1 Clinical trials development	3
1.2 Specific aspects of phase I dose-finding clinical trials in oncology . .	4
1.3 Single-agent dose-finding clinical trials	8
1.4 Combination dose-finding clinical trials	11
1.5 Review of drug combinations phase I trials in oncology	13
2 State of the art: Dose-finding designs for single-agent trials in oncology	31
2.1 Standard designs for single-agent dealing with toxicity	31
2.1.1 Algorithm-based designs	31
2.1.1.1 Traditional “3+3” design	32
2.1.1.2 “A+B” scheme without dose de-escalation	33
2.1.1.3 Accelerated titration designs	34
2.1.2 Model-based designs	34
2.1.2.1 Continual Reassessment Method	34
2.1.2.2 Escalation With Overdose Control	44
2.2 Designs for single-agent trials dealing with both toxicity and efficacy	45
3 State of the art for combination designs: Competing designs for drug combination in phase I dose-finding clinical trials	47
4 A Bayesian dose-finding design for drug combination clinical trials based on the logistic model	71
5 A Bayesian Dose-finding Design for Clinical Trials Combining a Cytotoxic Agent with a Molecularly Targeted Agent	85

6 Phase I/II Dose-Finding Design for Molecularly Targeted Agent: Plateau Determination using Adaptive Randomization	102
7 Conclusion and discussion	122
Perspectives	127
A new Bayesian dose-finding method for an ordinal efficacy endpoint and a binary toxicity endpoint	129
Phase I/II Dose-Finding Design for Molecularly Targeted Agent with Continuous Efficacy Outcome	129
A Appendices	131
A.1 Lee and Cheung's skeleton	132
A.2 Clopper-Pearson Confidence Interval	136
A.3 Adaptive Rejection Metropolis Sampling within Gibbs Sampling and reminders	138
A.4 Isotonic regression and Pool-Adjacent-Violators Algorithm	146
A.5 Beta regression	149
A.6 Posterior Predictive Loss	151
Bibliography	152
Résumé long	160
Abstract	169

List of Tables

1.1	Characteristics of the 162 drug combinations phase I trials reviewed.	26
2.1	Updated toxicity probabilities of the five dose levels after each newly included cohort using the CRM in the context of a real clinical trial.	43
3.1	Toxicity scenarios for the two-agent combinations.	55
3.2	Comparison of all dose-finding designs in terms of percentage of correct selection, percentage of patients allocated at the true MTD(s) during the trial, and mean number of observed DLTs throughout the trial when the aim is to select only one MTD.	56
3.3	Comparison of AISO, TSTAT, I2D, BCOPULA, and BGUMBEL designs in terms of percentage of correct combination selection for each level of agent 2 when selecting multiple MTDs.	58
3.4	Working models retained to study their influence on PCS for I2D, BCOPULA and BGUMBEL. These working models were constructed according to Lee and Cheung.	62
3.5	PCSs for I2D, BCOPULA and BGUMBEL according to the three working models.	62
3.6	PSCs for the I2D design according to the presence or absence of the interaction term γ in the model.	63
3.7	PCSs for POCRM using working model 2 with an increase in the number of orderings from 3 to 200.	63
3.8	Percentage of correct selection of POCRM using five different working models.	65
3.9	Mean number of patients allocated to each combination all along the trial ($N = 60$) for AISO, TSTAT and I2D designs.	67
3.10	Mean number of patients allocated to each combination all along the trial ($N = 60$) for POCRM, BCOPULA and BGUMBEL designs.	68
3.11	Percentage of combination recommendation for multiple MTDs selection using decision rules described in simulation section.	69
3.12	Percentage of combination recommendation for multiple MTDs selection using decision rules described in simulation section.	70
4.1	Toxicity scenarios for the two-agent combinations.	78
4.2	Percentage of correct selection in comparison of six dose-finding designs. Percentage of patients allocated to a true MTD during the trial and mean number of observed DLTs throughout the trial. . . .	79

4.3	Combination selection percentages of the proposed method for combination evaluation.	80
4.4	Additional toxicity scenario for the two-agent combinations.	80
4.5	Percentage of correct selection (PCS) and percentage of stopping for toxicity without MTD recommendation. Results of the sensitivity analysis with different prior distributions, using for β_0, β_3 normal distribution with high or low variability and for β_1, β_2 exponential or gamma priors with high or low variability.	81
4.6	Percentage of correct selection (PCS) and percentage of stopping for toxicity without MTD recommendation when toxicity is considered as a time-to-event outcome.	82
5.1	Prior distributions for model parameters.	92
5.2	Eight toxicity and efficacy scenarios for the combination of a cytotoxic agent (agent A) with an MTA (agent B).	96
5.3	Selection percentage of the ODC and EDC, the average number of patients treated at the ODC and EDC, and the average number of DLTs.	97
5.4	ODC (bold) and EDC (<i>italic</i>) selection percentages and the average number of patients treated at each combination.	97
5.5	Results of the sensitivity analysis.	98
5.6	Summarized data from the case-study combination clinical trial involving imatinib with paclitaxel.	99
6.1	Selection percentage and the number of patients (shown in parentheses) allocated to each dose under the TC design and the proposed MTA-RA and MTA-PM designs. In bold are given the optimal doses, and “correct” doses are underlined.	116
6.2	Selection percentages of the optimal dose and “correct” doses without random effect.	117
6.3	Sensitivity analysis of the proposed MTA-RA and MTA-PM designs. In bold are given the optimal doses.	118
6.4	Selection percentages of the optimal dose and “correct” doses using random effect model on generated data with correlatio term δ	120
6.5	Sensitivity analysis with different link functions and prior distributions.	121

List of Figures

1.1	MTD definition.	5
1.2	Extracted page of CTCAE v4.0.	7
1.3	Source from BioMedTracker, 2012. ©	9
1.4	Optimal dose definition.	11
1.5	Agent levels.	12
1.6	Partial known orderings for the combination of two cytotoxic agents.	12
1.7	Two possible orderings satisfying the partial order.	13
1.8	(a) Partial known ordering between combinations. (b) Possible orderings between combinations according to increasing DLT rates.	24
1.9	Flow chart of the publications found from the Medline search.	24
1.10	Control quality of the 162 trials reviewed.	25
2.1	“3+3” strict traditional escalation rule.	32
2.2	"A+B" escalation design.	33
2.3	Continual Reassessment Method (CRM).	35
2.4	(a) Scenarios and (b-f) Comparison of percentage of selections for the “3+3” and the CRM. The true MTD is given in red.	39
2.5	Probability tree for dose recommendation of z additional patients.	41
3.1	(a) Partial order relationships between combinations. (b) Three possible orderings consistent with the partial order retained for POCRM among all possible orderings for the simulation study.	51
3.2	Convergence curves for scenarios 2–5.	57
3.3	Ordering causing a fall in PCS in scenario 8 with random orderings for POCRM.	63
3.4	Five different working models used in the simulation study for POCRM.	64
4.1	Posterior densities of the toxicity probability for each combination. The dashed line represents the toxicity target.	75
4.2	Probability of achieving the targeted interval, $P(\pi_{j,k} \in [\theta - \delta; \theta + \delta])$ at the end of the trial for each combination. The dashed vertical line corresponds to the highest probability, which determines the combination selected as the MTD ((2, 3) in this example).	76
5.1	Illustration of combination zones for the start-up phase.	94

6.1	Two possible plateau estimations depending on the highest dose explored (not only) when extended our design to unimodal relationships.	104
6.2	Illustration of dose-toxicity and dose-efficacy relationships for MTAs.	119
A.1	Illustration of the construction of H intervals.	132
A.2	The idea of Lee and Cheung's skeleton.	135
A.3	Construction of the function envelop for ARS.	139
A.4	Evolution of the envelop after several rejection.	141
A.5	Adaptive rejection function for ARMS.	144
A.6	Various shapes of beta densities.	149
A.7	Détermination de la dose maximale tolérée.	161
A.8	Détermination de la dose optimale.	162
A.9	Relation d'ordre partielle entre les combinaisons de deux agents cytotoxiques.	164

Scientific production

Accepted papers:

- **Riviere MK**, Dubois F, Zohar S. Competing designs for drug combination in phase I dose-finding clinical trials, *Statistics in medicine* 2014.
- **Riviere MK**, Dubois F, Zohar S. SIM-12-0373 and sim-14-0573: Response to 'Comments on Competing designs for drug combination in phase I dose-finding clinical trials' by G. Yin, R. Lin and N. Wages, accepted in *Statistics in medicine*.
- **Riviere MK**, Yuan Y, Dubois F, Zohar S. A Bayesian Dose-finding Design for Clinical Trials Combining a Cytotoxic Agent with a Molecularly Targeted Agent, *Journal of the Royal Statistical Society - Series C* 2014.
- **Riviere MK**, Yuan Y, Dubois F, Zohar S. A Bayesian dose-finding design for drug combination clinical trials based on the logistic model, *Pharm Stat.* 2014, 13(4):247-57.
- **Riviere MK**, Le Tourneau C, Paoletti X, Dubois F, Zohar S. Designs of drug combinations phase I trials in oncology: a systematic review of the literature, accepted in *Annals of Oncology*.

Papers in revision:

- **Riviere MK**, Dubois F, Jourdan JH, Zohar S. Phase I/II Dose-Finding Design for Molecularly Targeted Agent: Plateau Determination using Adaptive Randomization, in review in *Statistics in medicine*.

Paper in progress:

- A Bayesian dose-finding method for an ordinal efficacy endpoint and a binary toxicity endpoint, in collaboration with M. Ezzalfani and S. Zohar.

Conferences:

- 2013:
 - ISCB - International Society for Clinical Biostatistics, Munich.
Oral presentation “A Bayesian Dose-finding Design for Clinical Trials Combining a Cytotoxic Agent with a Molecularly Targeted Agent”.
 - International Meeting - Statistical Methods in Biopharmacy, Paris.
Oral presentation “Phase I/II Dose-Finding Design for Molecularly Targeted Agent: Plateau Determination using Adaptive Randomization”
- 2012: Days of Biopharmacy and health (“Journées Biopharmacie et Santé”), Paris.
Oral presentation “A Bayesian dose-finding design for drug combination clinical trials based on the logistic model”.
- 2011: Days of Biopharmacy and health (“Journées Biopharmacie et Santé”), Paris.
Oral presentation “Competing designs for drug combination in phase I dose-finding clinical trials”.

National and International Collaborations:

- Y. Yuan, MD Anderson Cancer Center, Houston, Texas, USA.
- C. Letourneau et X. Paoletti, Curie Institute, Paris, France.

Grants:

- Grant CIFRE number 2011/0900.
- Grant “Aide à la mobilité internationale 2012” from Paris 7 Denis Diderot University.

Glossary

- AE = Adverse Event.
- ANRT = Association Nationale de la Recherche et de la Technologie
- ARMS = Adaptive Rejection Metropolis Sampling.
- CIFRE = Conventions Industrielles de Formation par la REcherche.
- CRM = Continual Reassessment Method.
- CTCAE = Common Terminology Criteria for Adverse Events = list of adverse events subdivided into organ/symptom categories that can be related to the anticancer treatment and that enable to categorize AEs into grades.
- DLT = Dose Limiting Toxicity = Toxicity: toxicity of grade 3 or more, where it is considered that the patient had a serious event.
- EMA = European Medicines Agency.
- EWOC = Escalation With Overdose Control.
- FDA = Food and Drug Administration.
- GS = Gibbs Sampling.
- INSERM = Institut National de la Santé Et de la Recherche Médicale (National Institute of Health and Medical Research).
- IRIS = Institut de Recherches Internationales Servier
- LD10= Dose that is Lethal in 10% of the pre-clinical treated mice .
- MCMC = Monte Carlo Markov Chains.
- MED = Minimum Effective Dose.
- MRI = Magnetic Resonance Imaging
- MTA = Molecularly Targeted Agent.

- MTD = Maximum Tolerated Dose: highest dose that can be administered to a patient without too much adverse events, dose closest to a predetermine target toxicity.
- PCS = Percentage of Correct Selection.
- PET-Scan = Positron Emission Tomography
- TDL = lowest dose that produces side effects and for which Twice that Dose is not Lethal.
- RECIST = Response Evaluation Criteria In Solid Tumors.
- RP2D = Recommended Phase 2 Dose.

CIFRE PhD

My PhD is the result of a public-private partnership. Through ANRT (“Association Nationale de la Recherche et de la Technologie”), which awards grants to companies that hire a PhD student for three years to set up a research collaboration with a public laboratory, I obtained a CIFRE grant (number 2011/0900). My PhD is performed with the pharmaceutical company IRIS (Institut de Recherches Internationales Servier), second pharmaceutical French global group. I am working in the division Methodology and Data Valorisation, directed by Maylis Coste, in the oncology department, directed by Christine Gabarroca, and recently also by Frédéric Dubois, who is my scientific supervisor. My work is guided by the research laboratory, INSERM (Institut National de la Santé Et de la Recherche Médicale), Unit 1138 based at “Centre de Recherches des Cordeliers”, team 22 which specializes in information sciences to support personalized medicine. I was directed by Sarah Zohar, PhD, an expert in early phase dose-finding.

PhD Objectives and scientific progress

My PhD dissertation deals with dose-finding phase I clinical trials in oncology. Our objective is to develop innovative Bayesian adaptive designs in phase I. My work first focuses on designs involving only toxicity. In this context, we decided to concentrate on combination designs, as in current practice combination therapies are of major interest. More and more drug combination trials in phase I have been introduced and methodological research is recent, with a few designs developed in the last 10 years. We compare existing designs in the literature via an extensive simulation study, and then propose our own approach for combination designs dealing with toxicity. Then, as we were involved in real clinical development within the pharmaceutical company IRIS, we became aware that new types of molecules, molecularly targeted agents (MTA), have emerged as an alternative or a complement to cytotoxic agents. These molecules raised different assumptions than with cytotoxic agents, and new designs were needed that also involve the modelization of efficacy. Therefore, another part of my PhD research was dedicated to the development of phase I designs for MTA involving both toxicity and efficacy in single-agent and combination trials.

Manuscript organization

For the dissertation, we have decided to organize the manuscript differently than the chronological order. We have chosen to separate the single-agent and combination methodology. In a first chapter, we introduce the concept of dose-finding in phase I oncology and detail the specificities of MTA and combination trials. In a second chapter, we then present the existing methods for single-agent phase I

trials (with toxicity alone or both toxicity and efficacy). Then, in a third chapter, we compare the existing designs for combination phase I trials that was recently published. We propose a Bayesian adaptive design for MTA in single-agent phase I clinical trials in chapter 4. Chapters 5 and 6 introduce the designs we developed for combination studies dealing with first cytotoxic agents, and then cytotoxic agent combined with a MTA. Finally, we conclude this manuscript with a discussion and perspectives.

Chapter 1

Introduction

1.1 Clinical trials development

The development of a new drug for a given therapeutic indication is carried out in four or five clinical trial phases preceded by a pre-clinical phase.

The pre-clinical phase consists of the study of the molecule, its structure, and its effect on cells. The new drug is administered to laboratory animals, such as mice, dogs or monkeys, in order to determine the effect on the animal's behavioral and biological level. In this pre-clinical phase, safety data are collected in order to assess the tolerability of the new drug. Pharmacokinetics/pharmacodynamics data are also gathered in order to study the becoming of the active substance after the administration in the body, and characterize the new molecule. Sometimes efficacy data such as the ability to target some pathways are also of main interest. The first dose level that will be used in humans is determined from the observed pre-clinical data. Its definition depends on the disease and its life-threatening character.

Clinical trials involving new drugs are commonly classified into five phases. Each phase of the drug approval process is treated as a separate clinical trial. The drug-development process normally proceed through all four phases over many years. If the drug successfully passes through phases 0, I, II, and III, it will usually be approved by the regulatory authority for use in the general population.

Sometimes, a phase 0 is performed before phase I. Phase 0 trials consist in administering a very small dose of treatment for a short period of time on a small group of patients (about 10-20). These trials are used to obtain information on how the human body reacts to the drug and study its possible effects. They enable verification of whether a drug exerts the expected action or effect. The main difference between phase 0 and other phases is that participants cannot expect any immediate and direct benefit from participating in this exploratory trial. However, as the doses are very low, there are also less risks than in phase I. Phase I trials are the first stage of testing in human subjects. The goal of phase I trial is to evaluate the safety (and feasibility) of the treatment and identify its side effects.

They enable investigators to define and characterize new treatments in humans to set the basis for later investigations of efficacy. For non-life-threatening diseases, phase I trials are usually conducted on human volunteers. In life-threatening diseases such as cancer or AIDS, phase I studies are conducted with patients because of the aggressiveness and possible harmfulness of the treatments, possible systemic treatment effects, and the high interest in the new drug's efficacy in those patients directly. Once the initial safety of the study drug has been confirmed in phase I trials, phase II trials are performed on larger groups and are designed to establish the efficacy of the drug and confirm the safety identify in phase I. It is the first evaluation of efficacy. Phase III studies compare the efficacy of the drug, in comparison with current "gold standard" treatment or a placebo. They are usually randomized controlled multi-center trials on large patient groups. Phase IV trials, or pharmacovigilance, involve the safety surveillance and ongoing support of a drug after it receives permission to be sold. Phase IV consists of the collection, detection, assessment, monitoring, and prevention of adverse effects with pharmaceutical products under prescription. Pharmacovigilance enables to detect rare or late adverse events (AEs) in the general population that was not previously detected in clinical trials.

1.2 Specific aspects of phase I dose-finding clinical trials in oncology

Phase I trials are the first stage of human experimentation with a new drug or combination [16]. They often involve drugs that have been tested extensively in the laboratory and on animals with encouraging results, but have not yet been given to humans. Phase I trials involve ethical concerns. Indeed to be safe, the treatment starts at low doses but they are probably not high enough to elicit a beneficial effect. But if patients are exposed to doses with a higher toxicity rate, the drug can be very harmful. The goal in dose-finding is to obtain enough information on toxicity in order to estimate correctly the most appropriate dose with the fewest patients.

Only a small number of participants are enrolled in phase I trials, usually 15 to 50 patients. Trial participants are divided into small groups, known as cohorts. At this early stage of development, as efficacious and safe dosing is unknown, and information is available at best from pre-clinical studies. Therefore, phase I trials are sequential dose-escalation procedures: treatment generally begins at a low dose level very likely to be safe (starting dose), and then small cohorts of patients are treated at progressively higher doses (dose escalation) until the drug-related toxicity reaches a predetermined level. Patients in phase I trials are sometimes the first to try the new cancer drug. Phase I clinical trials in oncology include patients who failed standard treatments or are in the last stages of the disease; the new drug may be the last remaining chance for effective treatment. In these

phases, involving very ill patients with advanced cancer often highly pre-treated, the response rate is very low, but patient benefit or anti-tumor activity is not the goal of these studies.

The primary aim of phase I clinical trials is to determine the maximum tolerated dose (MTD). Following Storer [67], the statistical formulation of the problem is to select a dose level from several available doses, with a toxicity probability closest to a given target [54, 23, 43]. For example in Figure 1.1, the target toxicity probability is fixed to 0.30, and the dose level closest to the posterior estimated toxicity probabilities is dose level 7. The MTD can also be defined as the highest dose level with acceptable toxicity rate.

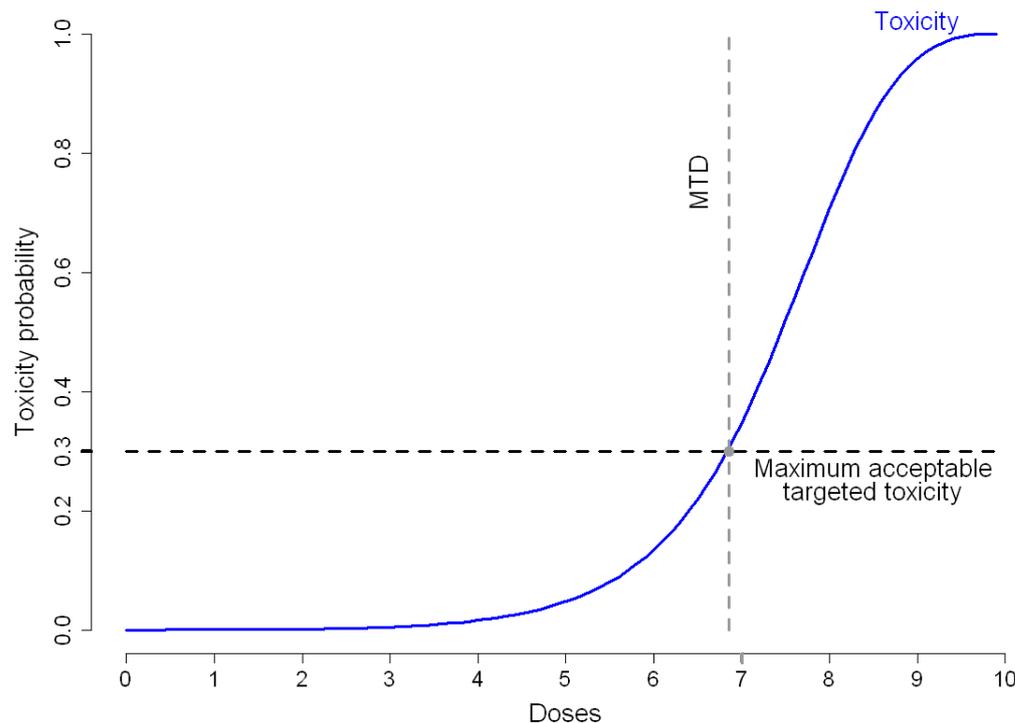


Figure 1.1: MTD definition.

Drugs used in phase I trials may have been developed by pharmaceutical or biotechnology companies. Some phase I trials test new uses for drugs that have already been approved by the Food and Drug Administration (FDA) in the United States or by the European Medicines Agency (EMA) in Europe, or test drugs given for the first time in combination. Drugs may also be administered in different ways, such as by regional therapy.

The first dose level in phase I clinical trials in oncology is in general defined as 1/10 of the dose lethal in 10% of the mice treated in pre-clinical trial (LD10) or as 1/3 of the lowest dose that produces side effects in dogs and for which twice that dose is not lethal (TDL). The choice of subsequent dose levels is determined

from either animal experimentation, or a Fibonacci mathematical sequence (the next dose level in the sequence is equal to the sum of its two predecessors), or by physicians.

At the beginning of anti-tumor drugs development and until recently, the main agents developed were cytotoxic agents. Cytotoxic agents are chemotherapy drugs that treat malignancies by directly killing tumor cells that divide rapidly. The main assumptions for phase I using cytotoxic drugs is that both the dose-toxicity and the dose-efficacy relationships are monotonic and increasing with the dose level [13]. Therefore, the higher the dose, the higher the efficacy, and so our ability to kill cancer is higher, but also the toxicity and the severity of the observed adverse events are higher. Thus, with very high doses of these therapies, physicians may be able technically to eliminate the cancer completely, but it would result in the patients' deaths and therefore would be of no interest. That is why the dose to select in phase I oncology trials is a compromise between toxicity and efficacy. Because of the paradigm "More is better" [34], only the toxicity need to be studied, and the aim is to select the highest dose level under toxicity restrictions in order to select an efficient dose enough but that would be tolerable for patients. That is the notion of MTD defined previously.

The notion of MTD is defined in terms of observed toxicity data of the patients treated using the notion of dose limiting toxicity (DLT) under valid toxicity criteria. Drug toxicity is considered tolerable if the toxicity is acceptable, manageable, and reversible. Drug safety has been standardized for oncological studies by the establishment of the Common Terminology Criteria for Adverse Events (CTCAE), of the U.S. National Cancer Institute. This is a very large list of adverse events subdivided into organ/symptom categories that can be related to the anticancer treatment. Each AE can be categorized into a grade:

- grade 0, No AE or normal
- grade 1, Mildly (elevated/reduced)
- grade 2, Moderate
- grade 3, Serious/severe
- grade 4, Very serious or life-threatening
- grade 5, Fatal

The most recent version of CTCAE is 4.0, and includes many pages of symptoms description classified into grades. As an example, Figure 1.2 shows an extracted page of CTCAE.

Gastrointestinal disorders					
Adverse Event	Grade				
	1	2	3	4	5
Colonic fistula	Asymptomatic; clinical or diagnostic observations only; intervention not indicated	Symptomatic; altered GI function	Severely altered GI function; bowel rest, TPN or hospitalization indicated; elective operative intervention indicated	Life-threatening consequences; urgent intervention indicated	Death
Definition: A disorder characterized by an abnormal communication between the large intestine and another organ or anatomic site.					
Colonic hemorrhage	Mild; intervention not indicated	Moderate symptoms; medical intervention or minor cauterization indicated	Transfusion, radiologic, endoscopic, or elective operative intervention indicated	Life-threatening consequences; urgent intervention indicated	Death
Definition: A disorder characterized by bleeding from the colon.					
Colonic obstruction	Asymptomatic; clinical or diagnostic observations only; intervention not indicated	Symptomatic; altered GI function	Hospitalization indicated; elective operative intervention indicated; disabling	Life-threatening consequences; urgent operative intervention indicated	Death
Definition: A disorder characterized by blockage of the normal flow of the intestinal contents in the colon.					
Colonic perforation	-	Symptomatic; medical intervention indicated	Severe symptoms; elective operative intervention indicated	Life-threatening consequences; urgent intervention indicated	Death
Definition: A disorder characterized by a rupture in the colonic wall.					

Figure 1.2: Extracted page of CTCAE v4.0.

These AEs are generally considered as “a necessary evil” regarding the expected benefit of cytotoxic treatments. There is no universally accepted definition of DLT. Usually, a toxicity of grade greater than or equal to 3 is considered as dose limiting. Nevertheless, for example, grade 2 AEs with a long duration can also be considered as a DLT. The definition of DLT for a specific clinical trial has to be defined explicitly by the physicians in advance of the beginning trial. The toxicity data is a binary outcome equal to 0 (no toxicity) if no DLT was observed and 1 (toxicity) if at least a DLT was observed. Generally, in phase I cancer studies, the DLTs are evaluated only for a short period of time at the beginning of the treatment. Classically DLTs are based on the first cycle of treatment (cycle 1 toxicity). The duration of one cycle of treatment can vary but is usually 3 weeks. However, patients continue to be treated on more cycles (sometimes 8 cycles depending on whether the progression of the disease was observed, the occurring DLTs, or patients’ health status if they are still alive). This means that all the toxicity outcomes that occur after the first cycle of treatment are not taken into account into the statistical model used to determine the MTD, which is then the dose recommended for phase II (recommended phase 2 dose, RP2D). In practice, for safety and ethical reasons, it is not possible not to consider DLTs occurring in further cycles of treatment and therefore physicians have to arbitrarily adjust the dose level using their own experience. Due to these “late toxicities” or to other medical considerations, the final RP2D retained at the end of the study can be different from the dose level that would be recommended by the model.

To summarize, in phase I oncology trials, the information available is restricted and consists of binary toxicity data evaluated on a short period with a small sample size of advanced cancer patients.

1.3 Single-agent dose-finding clinical trials

Phase I cytotoxic clinical trials in oncology involve several ethical concerns. Indeed, in order to gather information about the dose-toxicity relationship, it is not possible to include a large number of patients and randomize them at each different dose level considered in the trial. As previously stated, patients cannot be treated with dose levels that would be greater than the MTD, as they would be exposed to very high toxicity. Moreover, the number of patients treated at low dose levels should be minimized, as for cytotoxic agents efficacy is assumed to accompany with toxicity, therefore low dose levels without AEs are considered inefficient. In addition, as the total sample size is very limited, confidence intervals are large. For logistical constraints, costs constraints, recruitment difficulties, and exposure of patients to non-optimal doses, the sample size can often not be increased. Nevertheless, a high percentage of mis-identification of the RP2D has been observed, leading to high failure of clinical trials in further phases. In fact, the success rate in phase II and III is globally lower in oncology than in other therapeutic areas according to a recent paper in *Nature* [2] and to BioMedTracker (see Figure 1.3).

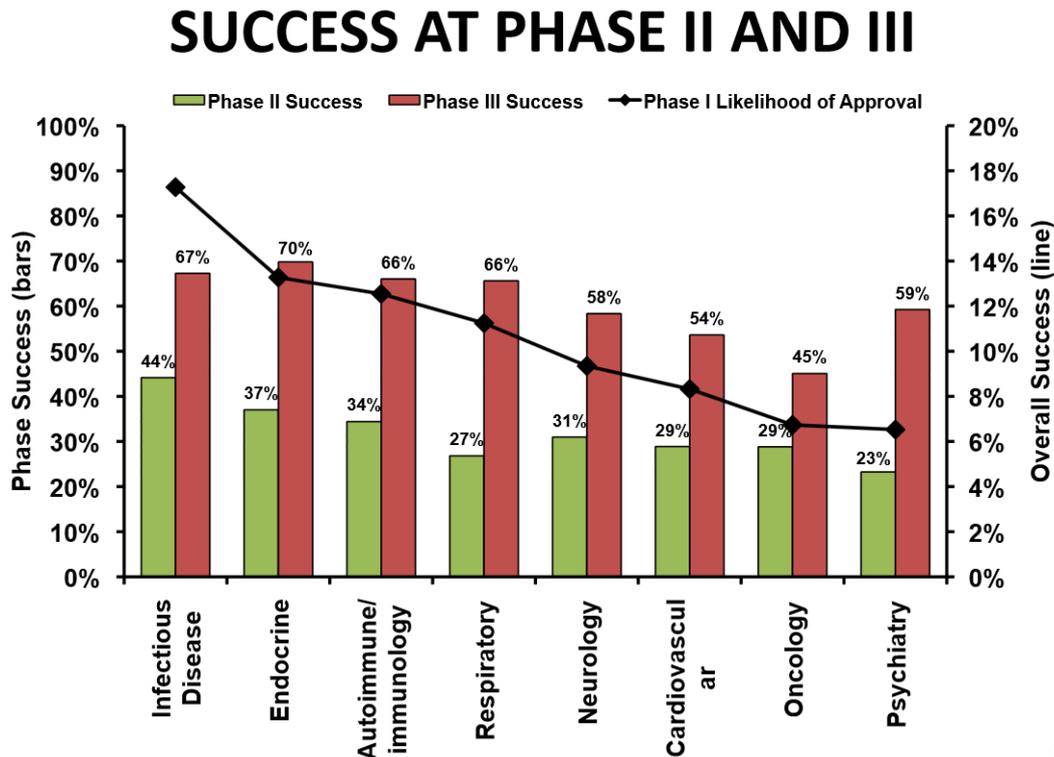


Figure 1.3: Source from BioMedTracker, 2012. ©

The failure rate in phase II is 71% and 55% in phase III, which is very high. A part of this failure is not necessarily due to the discovery of an ineffective anti-tumor agent but also to a mis-identification of the appropriate dose in phase I. Therefore, more resources should be placed in phase I. Indeed, as the sample size is much higher in further phases, when the RP2D found at the end of phase I is far the (“best”) dose that should have been recommended, more patients are exposed to wrong dose level in phases II and III raising also ethical concerns for a larger number of patients. Therefore, there has been many discussion and compromise regarding the sample size in phase I, but statisticians still need to convince physicians to have adequate sample size.

As the total number of patients is small and all dose levels cannot be explored, the models proposed by statisticians in phase I are limited. It is not feasible to have toxicity probability estimations that are reliable for each dose level. Therefore, the aim is not to capture the entire dose-toxicity relationship, but to correctly estimate locally around the MTD the toxicity probabilities in order to recommend the appropriate dose level at the end of the trial. In this case, toxicity rates corresponding to dose levels far from the MTD would not be properly estimated, but the goal is not to estimate all toxicity rates but to recommend a suitable dose level for phase II.

Several issues are also raised for physicians and statisticians before the trial onset [92]. Which dose range should be chosen (number, spacing between dose levels, ...)? Which starting dose? How many subjects should be enrolled? etc... Some clinical choices are of the responsibility of physicians, but both physicians and statisticians are working together to set up the study.

From the end of World War II, the main oncology agents developed were cytotoxic. Nevertheless, in 2000 the first MTA, trastuzumab (Herceptin®), was developed by Roche in the breast cancer. MTAs have emerged as a new treatment option in oncology that have changed the practice of cancer patient care [44, 46, 45, 62]. These type of agents have a different mechanism of action compared to cytotoxic agents; instead of killing cells, they target specific pathways. For example, the targeted agent trastuzumab acts by linking the receptor HER2 to stop the growth of cancer cells. This MTA is only efficient on a targeted subgroup of the population, women with the biomarker HER2+; that is, if cells have too much of the protein HER2, trastuzumab is efficient and increases overall survival, whereas for women HER2-, that is if the cancer cells have normal amounts of HER2, the drug is ineffective. Since 2000, many MTAs were developed in several cancer indications with approximately one new MTA every year.

For cytotoxic agents, it is assumed that the more toxic the drug, the more efficient it is. However, for many MTAs this assumption is not satisfied. For some MTAs, it has been observed the occurrence of a plateau of efficacy when increasing the dose [33]. Indeed, for example, when all the receptors targeted are already linked, an increase of dose level would result in a saturation of receptors in the body, and therefore a plateau of efficacy can be observed. In this case, it is more appropriate to identify the dose at the breaking point of the plateau while maintaining toxicity restrictions rather than the highest possible dose under toxicity constraints. We call this dose the “optimal dose”. In Figure 1.4, for example, the optimal dose level is dose 3 whereas the MTD is dose 7. Indeed, there is no need to increase the dose, and so the toxicity for the same efficacy. All dose levels between the optimal dose and the MTD would be “correct” to administer, as they have the highest possible efficacy under toxicity constraints, but it is more ethical to recommend the optimal dose, because for the same efficacy level, this dose has the lowest toxicity.

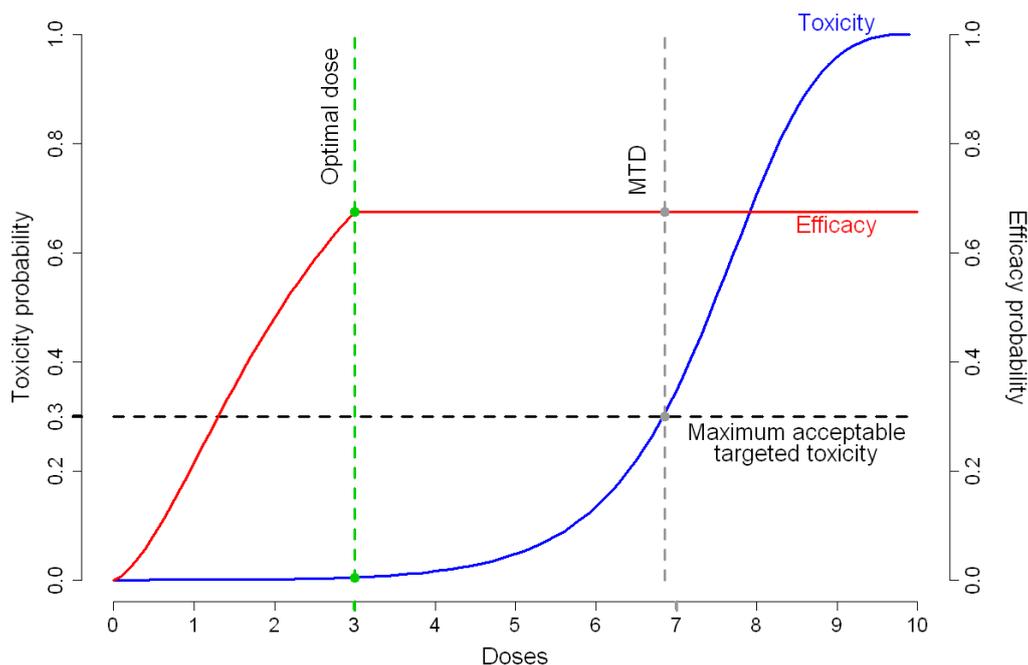


Figure 1.4: Optimal dose definition.

Other phenomena can be observed with the MTA, as for instance an increase and then decrease of efficacy with the dose. In this case, the aim is to recommend the dose level at the mode of the dose-efficacy curve.

1.4 Combination dose-finding clinical trials

Most statistical model-based or algorithm-based methods have been developed for cytotoxic single-agent phase I dose-finding clinical trials [13]. In this context, it is assumed that the toxicity of a single-agent is monotonic and increases with the dose, as does the efficacy.

With recent progress in the field of oncology, it is rare to find new molecules that outperform already existing therapeutic strategies. Moreover, cancers can develop diverse mechanisms of resistance to therapy with single agents. That is why, in different areas, but especially in cancer studies, more and more combination studies are introduced [3]. By combining two or more agents, investigators wish to increase the overall anti-tumoral action and survival due to a potential synergistic effect between drugs in terms of efficacy. As a result, it is difficult to suppose that each molecule will act independently in terms of toxicity. We suppose that columns represent levels of agent 1 and rows levels of agent 2, with the increasing way defined as in the figure below:

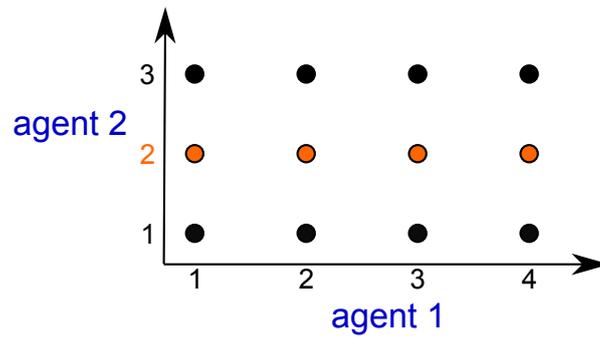


Figure 1.5: Agent levels.

In dose-finding studies, physicians wish to gradually increase toxicity during the dose escalation procedure. However, when combining several agents, the ordering of toxicity probabilities is not completely known. For instance, the combination of two cytotoxics, for which toxicity is increasing with the dose level of both agent, can induce an ordered subset of toxicity [15, 77, 78] (Figure 1.6). When we fix one of the agents, toxicity increases when the other agent is increased, leading to a partial order between combinations. The symbol “<” represents an order relationship between the toxicity probabilities describing the statement “is inferior to”.

$$\begin{array}{cccccc}
 D_{1,3} & < & D_{2,3} & < & D_{3,3} & < & D_{4,3} & < & D_{5,3} \\
 \vee & & \vee & & \vee & & \vee & & \vee \\
 D_{1,2} & < & D_{2,2} & < & D_{3,2} & < & D_{4,2} & < & D_{5,2} \\
 \vee & & \vee & & \vee & & \vee & & \vee \\
 D_{1,1} & < & D_{2,1} & < & D_{3,1} & < & D_{4,1} & < & D_{5,1}
 \end{array}$$

Figure 1.6: Partial known orderings for the combination of two cytotoxic agents.

In contrast to single-agent treatments, the underlying assumption for the agents do not enable the determination of a full ordering among combinations. Even when a partial ordering is known, it is still difficult to decide how to escalate or de-escalate a combination of doses. Indeed, on a diagonal, there is no knowledge about which combination is more toxic. For instance, in Figure 1.6, it is not known which one of $D_{1,2}$ or $D_{2,1}$ is more toxic prior to the trial. Several orderings satisfying the partial orders are possible, as for example in Figure 1.7. Therefore, it is senseless to use single-agent dose-finding methods for combination studies.

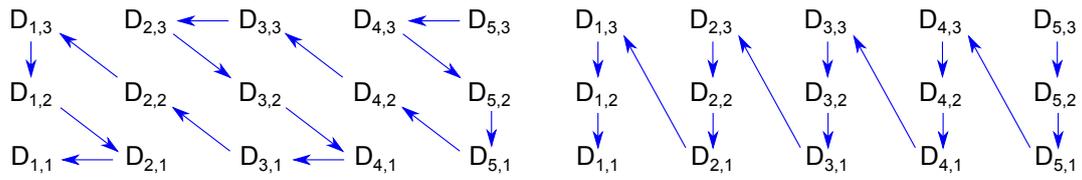


Figure 1.7: Two possible orderings satisfying the partial order.

1.5 Review of drug combinations phase I trials in oncology

We performed a systematic review of the literature of all drug combinations phase I trials published the last three years between 1 January 2011 and 31 December 2013, where at least two drugs were planned to undergo dose escalation. Our aim was to determine what were the current practice in combinations trials. This review was the object of a paper that is in press in *Annals of Oncology*.

Review

Designs of drug-combination phase I trials in oncology: a systematic review of the literature

M-K. Riviere,^{1,2} C. Le Tourneau,^{3,4} X. Paoletti,³ F. Dubois,² S. Zohar¹

1: INSERM, U1138, Team 22, Centre de Recherche des Cordeliers, Université Paris 5, Université Paris 6, Paris, France;

2: IRIS (Institut de Recherches Internationales Servier), Suresnes, France;

3: Institut Curie INSERM, U900, Paris, France ;

4: Department of Medical Oncology, Institut Curie, Paris, France

Corresponding author: Ms Marie-Karelle Riviere

INSERM UMRS 1138, Team 22, Centre de Recherche des Cordeliers

Escalier D, 1^{er} étage

15, rue de l'Ecole de Médecine

75006 Paris, France

Tel: +33-1-4427-7676

Email: marie-karelle.riviere@crc.jussieu.fr

This manuscript contains 12 pages, 1 table, 3 figures and 34 references

Word count: text = 3088 words (max = 4000), word count abstract = 260 words (max = 300)

Abstract

BACKGROUND

Combining several anticancer agents can increase the overall anti-tumor action, but at the same time, it can also increase the overall observed toxicity. Adaptive dose-escalation designs for drug combinations have recently emerged as an attractive alternative to algorithm-based designs, and they seem more effective in combination recommendations. These methods are not used in practice currently. Our aim is to describe international scientific practices in the setting of phase I drug combinations in oncology.

MATERIAL AND METHODS

A bibliometric study on phase I dose-finding combination trials was conducted using the MEDLINE® PubMed database between January 1, 2011 and December 31, 2013. Sorting by abstract, we selected all papers involving a minimum of two-agents and then retained only those in which at least two agents were dose-escalated.

RESULTS

Among the 847 references retrieved, 162 papers reported drug-combination phase I trials in which at least two agents were dose-escalated. In 88% of trials, a traditional or modified 3+3 dose escalation design was used. All except one trial used a design developed for single-agent evaluation. Our study suggests that drug-combination phase I trials in oncology are very safe, as revealed by the calculated median DLT rate of 6% at the recommended dose, which is far below the target rate in these trials (33%).

We also examined requirements of phase I clinical trials in oncology with drug combinations and the potential advantages of novel approaches in early phases.

CONCLUSION

Efforts to promote novel and innovative approaches among statisticians and clinicians appear valuable. Adaptive designs have an important role to play in early phase development.

Keywords: Drug combinations, Phase I trials, Dose-finding

Introduction

Phase I trials in oncology are dose-finding studies that seek to determine the dose to recommend for further evaluation (recommended phase II dose [RP2D]). These trials are designed to obtain reliable information on the safety, pharmacokinetics, and mechanism of action of a drug. In oncology, dose-finding studies focus on determining the highest dose of a new drug with acceptable toxicity [1-2]. They are subject to the ethical constraint of minimizing the number of subjects treated at unacceptable toxic dose levels. Toxicity is measured as a binary endpoint, denoted as “dose limiting toxicity” (DLT), mainly using National Cancer Institute Common Toxicity Criteria. Most dose-escalation methods were developed for cytotoxic agents with the assumption that toxicity increases with dose in a monotonic fashion. Therefore, the RP2D has traditionally been the highest safe dose, called the “maximum tolerated dose” (MTD). These methods were specifically designed for the evaluation of single agents. In clinical practice, the traditional “3+3” dose escalation design or a modification thereof are the most frequently used dose-escalation methods in phase I trials [3].

Drug combinations have been introduced with the goal of improving treatment efficacy by increasing overall anti-tumor activity and, presumably, survival. Successful drug combinations include a combination of cytotoxic agents for the treatment of germ-cell tumors and lymphoma, [polychemotherapy for the treatment of germ-cell tumors \[4-5\]](#), [combinations of trastuzumab with a taxane for HER2-positive breast cancer \[6\]](#), and [a combination of BRAF and MEK inhibitors \[7\]](#). Although it can reasonably be assumed that toxicity increases with dose for a single drug, the determination of the relationship between toxicity and doses of multiple drugs remains elusive. When combining two or more agents, there may be a potential synergistic effect, not only in terms of efficacy, but also in terms of toxicity [8]. Therefore, when combining several agents, the ordering between combinations according to their DLT rates is important. However, only partial ordering of DLT rates can be anticipated when the dose of only one drug is being escalated, whereas the dose of the other drugs in the combination is kept fixed (Figure 1). [That is, referring to Figure 1a, in a row \(or column\), one agent is fixed while the other is increased. In this case, the DLT rates are increased with the dose of the agent. All these order relations in rows and columns \(shown with the symbol inferior “<”\) lead to “a partial ordering of DLT rates” given in Figure 1a.](#) For example, if 2 agents with 3 dose levels

are considered, when a monotonic and increasing relationship is assumed with respect to both agents then a partial toxicity order is known between the 9 combinations. The lowest combination is dose level 1 of agent 1 combined to dose level 1 of agent 2 (1,1) and the highest combination is dose level 3 of agent 3 combined with dose level 3 of agent 2 (3,3). Presumably, combination (1,2) is less toxic than (2,2), which is also presumably less toxic than (3,2), etc. **However, on a diagonal, when the dose of one agent is increased while the dose of the other is decreased, it is not known which combination is more toxic. For instance, is the (1,2) combination more or less toxic than the (2,1) combination?** Therefore, **several toxicity orderings between combinations are possible (two examples are given in Figures 1b and 1c).**

In practice, drug-combination phase I trials raise several challenging points to be defined prior to the trial onset [9-17]: 1) starting dose of each agent; 2) choice of the dose range of each agent and the number of combinations to be evaluated; and 3) total sample size that is strongly related to the number of possible combinations. In this study, we aimed at evaluating how drug-combination phase I trials in oncology have been designed in the last three years and what the principal investigator's choices were with regard to the dose range, number of combinations, and statistical design.

Material and methods

All drug-combination phase I trials published between January 1, 2011 and December 31, 2013 were reviewed (Figure 2). We restricted our review to phase I combination trials where at least two drugs were planned to undergo dose escalation. Trials involving radiation therapy or drugs other than cytotoxic agents and molecularly targeted agents (MTAs) were excluded. MTAs were defined in our review as anticancer agents that selectively target molecular pathways, as opposed to DNA, tubulin or cell division machinery. Hormonal therapy and biological therapeutics, such as immunotherapy, were included.

We performed a MEDLINE® PubMed search using the following terms: *"Clinical Trial, Phase I[ptyp] AND cancer[MeSH] AND "2011/01/01"[PDAT] : "2013/12/31"[PDAT AND (combination OR combine OR combined OR combining)]"*. Among 847 references retrieved, 162 papers reported on a drug combination phase I

trial meeting our inclusion criteria, 381 papers involved drug combinations where only one agent was dose-escalated, while the others were fixed (Figure 2).

The following data were recorded: the number of drugs undergoing dose escalation, the types of drugs (cytotoxic agent versus MTA), the number of dose levels planned for each drug, justification of the starting doses, number, choice and justification of drug combinations, dose-escalation design used, addition of drug combinations during the trial, number of patients included, and target toxicity level. We also performed a quality control analysis of the reviewed papers.

In this review, the lowest combination is defined as the combination corresponding to the lowest dose levels planned of each agent. A monotonic and increasing dose-toxicity relationship with respect to both agents signifies that when fixing one agent or the other to a certain dose independently, the DLT rate of the combination increases with the dose level of the remaining agent.

Results

Characteristics of the drugs

The 162 phase I trials involved 340 drugs that underwent dose escalation. In the majority of the trials, only 2 drugs underwent dose escalation (Table 1). Trials that involved only cytotoxic agents, only MTAs, and a combination of cytotoxic agents and MTAs were roughly equally distributed.

Dose levels

The median number of patients included per trial was 25 [range: 7-136] (Table 1). In 69% of cases, the starting combination in the trial was the one associated with the lowest dose level of each agent considered in the trial. The starting dose used in the trial was justified (short explanation or only references) in 35% of the trials, respectively (Figure 3). The dose levels of each agent involved in the combinations of the clinical trial were justified in only 47 publications (29%). [Results of a quality control analysis are provided in Figure 3.](#)

Dose combinations

The median number of planned combinations in the trial protocol was 5 [range: 2-16], 5.5 [range: 3-15] and 12 [range: 12-12] in trials combining two, three, and four drugs, respectively. The median number of actually evaluated combinations was 4 [range: 2-

12], 4 [range: 2-9] and 3 [range: 3-3] in trials combining two, three, and four drugs, respectively.

The median ratio of the number of planned combinations to the number of possible combinations (defined as the number of planned combinations divided by product of the number of doses levels of each agent) was 0.67 [range: 0.25-1], 0.24 [range: 0.17-0.63] and 0.13 [range: 0.13-0.13] in trials combining two, three, and four drugs, respectively.

Dose escalation method

In most trials, a traditional 3+3 or a modified 3+3 dose escalation design was used (Table 1). Only one trial used a design developed for combination trials. Most of the selected papers assumed a monotonic and increasing dose-toxicity relationship, in 62% of trials, whereas 38% of papers assumed only a partial monotonic and increasing dose-toxicity relationship.

In 24% of the trials, additional drug combinations were evaluated during the trial for safety reasons.

Safety

The DLT target rate associated with the recommended dose was 33% in most studies (Table 1). However, according to the number of patients and DLTs reported at the RP2D, the calculated median DLT rate at the recommended dose was 6% [range: 0%-40%]. Nevertheless, in only 4% of trials was the DTL rate estimated by the authors at the recommended combination for further studies.

In 3% of the studies, the trial was stopped at the first dose level due to DLTs. Five trials were stopped for reasons relating to over-toxicity; that is, the lowest combination evaluated in the trial was considered too toxic and the trial was abruptly halted without finding a tolerable combination. Fifty-six per cent of the trials found the MTD according to its initial definition, and 11% of trials found an MTD without observing any DLTs throughout the trial. In 48% of trials, the progress observed in the trial did not match the initial planned method. [The trials that did not match the intended plan were all 3+3 or modified 3+3 statistical designs. The main observed differences from planning were: \(1\) difference in the planned number of patients per cohort with no justification and \(2\) a different allocation rule during the trial.](#)

Discussion

Our study suggests that drug-combination phase I trials in oncology are safe. Overall, however, the starting doses of the drugs in the trials reviewed, as well as the dose levels and the dose-escalation steps, were barely justified. In addition, the dose levels explored in the drug-combination phase I trials included in our study did not reflect the entire space of possible drug combinations. In most of cases, dose levels seemed to be arbitrarily decided. It remains to be evaluated whether non-explored drug combinations would have been able to produce increased anti-tumor activity without jeopardizing patient safety.

Only a limited number of combinations were explored and only a sub-set of combinations was evaluated, despite the larger number of possible combinations. In our MEDLINE® PubMed search, the median ratio of the number of combinations considered to the number of possible combinations indicated that approximately one-third of the combinations were not considered for two-drug combinations. This indicates that trial investigators may have selected the combinations to be evaluated prior to the trial, and that some combinations were excluded without documented rationale. Exploring the entire combination space is obviously not feasible in practice. Nevertheless, the choice of the combinations to explore should not be limited by partial toxicity ordering. The design should have the possibility to explore any combination estimated to be the best. In fact, due to possible interactions between drugs, pre-selecting an arbitrary reduced sub-set of combinations induces a risk in selecting a combination with a DLT rate far from target toxicity. Even if the targeted DLT proportion was most often about 33% in the papers, the median DLT rate associated with the RP2D at the end of the trial was much lower. That could be a reason why an intermediate combination was added, in some cases, which induced a non-monotonic dose-toxicity relationship in some trials.

During the review of this paper, the question was raised whether the low DLT rate could be due to MTAs for which the toxicity profile is different. Indeed, for these non-cytotoxic agents, very low toxicities are often observed with sometimes cumulative low-grade toxicities that may become dose limiting. The cumulative low-grade toxicities partially explain deviance from the intended plan. An FDA guideline [18] reported: "...cancer vaccine trials have used the "3 + 3 design" and the results show that, except in very rare situations, an MTD for a cancer vaccine may not be identified. In these trials, the dose-toxicity curve may be so flat that the highest dose

that can be administered is limited by manufacturing or anatomic issues rather than toxicity. Therefore, this “3+3 design” may not be the most suitable approach to gathering information from early phase trials of cancer vaccines, and alternative trial designs should be considered.” They added that: “When no DLT is expected or achieved, optimization of other outcomes, such as the immune response, can be useful to identify doses for subsequent studies”.

For this reason, standard dose-finding designs dealing only with toxicity, such as the “3+3” do not seem appropriate for some biological agents [19]. First, it is true that the dose determination based on less than 33% DLT on the first cycle of treatment for molecularly targeted agents is problematic. These non-cytotoxic agents have different toxicity profiles than cytotoxic agents. One possible reason for the observed low DLT rate at the RP2D could be due to the DLT evaluation only on the first cycle of treatment. Physicians can observe no DLT on the first cycle but cumulative low-grade toxicities that become dose limiting with later cycles of treatment. For this reason, they decrease the recommended dose level for phase II (in contrast to the statistical design), rendering a low DLT rate (evaluated only on the first cycle) for this dose. All cumulative toxicity grades on all available cycles should be considered in the statistical analysis for dose recommendations. Furthermore, depending on the biological agent, several dose-efficacy relationships could be observed: (1) monotonic and increasing; (2) monotonic increasing and then reaching a plateau; and (3) monotonic increasing and then decreasing with the dose. In the latter two cases, only studying toxicity in the dose-finding process is not sufficient, and efficacy should also be considered. Therefore, alternative designs should be developed. Adapting the way of doing early phase clinical trials for these innovative molecules is important, but changing usual practices in oncology is very complex and difficult. If regulatory agencies were to give clear instructions, trial sponsors and investigators would need to apply them. There are published statistical designs proposing alternative methods [20-22], therefore statistics should not be a limited factor.

However, in calculating the median DLT rate for trials in which the combination involved cytotoxic agents, we observed a DLT rate of 4%. Therefore, we do not believe that this is due to the type of agent but rather to the use of the “3+3” algorithm, where the dose retained is the dose under 2 DLTs over 3 or 6 patients. Indeed, in the trials studied, either the combination level was associated with no (or very few) DLTs, or the highest dose level in the trial did not even reach the target

toxicity. Thus, due to the small number of patients (3 or 6) with the “3+3” design, the estimation is unreliable and often close to 0%. It should be noted that combination trials of MTAs included more patients at the RP2D than combination trials of cytotoxic agents, perhaps due to the uncertainty on overall toxicity discussed above. That could explain the small increased difference in DLT rates despite the toxicity profile of MTAs, as the estimation with a greater number of patients is more reliable.

Most of the drug-combination phase I trial designs used the traditional “3+3” design or a modified version. Recent dose-escalation designs have been developed for drug combinations but never employed in the trials reviewed [23-30]. In all but one trial reviewed, the dose-toxicity relationship was considered to be one-dimensional, whereas the reality involved several agents inducing a multi-dimensional issue. Most of the time, the problem was brought back into a one-dimensional space by pre-selecting combinations with a known toxicity order to be evaluated.

The methods for single agents do not always seem appropriate for combination phase I trials in which the doses of several drugs vary, as they are not designed to take a multi-dimensional space into account. Several alternative designs were proposed for either algorithm-based or design-based combinations that give the possibility to explore any appropriate combination in the entire combination space according to the accumulated data. Ivanova and Wang proposed an “up-and-down algorithm-type” method with isotonic regression [31] that was used recently in Gandhi et al. [32]. Conaway et al. developed a design for multiple agents based on partial orders [33] that was used in the publications reviewed in Jones et al. [34]. Other authors have proposed model-based designs in which the multi-dimensional feature of the entire combination space is taken into account. These methods allow considering the entire combination space that includes a large number of combinations with non-monotonic relationships. It should be noted that these methods do not permit exploring combinations that are estimated to be too toxic. In a recent comparison, based on simulations, Riviere et al showed that these designs were comparable and had high operational characteristics [30]. However, it is true that these designs have only been shown to be effective in simulation studies (Riviere et al, Stat Med 2014), and they require the involvement of a statistical expert.

In a recent editorial, Mandrekar [9] pointed out the importance of using adequate methods for the evaluation of combinations. Our bibliometric work supports

this editorial with a large and detailed study on clinical practice in the phase I settings for combination trials.

Our analysis did not include trials published in abstract form. Although this induces a selection bias, the present analysis still provides useful data that may help improve the design of drug-combination phase I trials.

In our review, we did not state that both agents must be administered at their single-agent MTD when in combination. As two agents can have a synergistic, antagonistic or independent effect on toxicity, the question of achieving doses (for each agent) that nearly approximate the recommended phase II dose is up for debate. It is a strong assumption that the addition of both agents at their MTD would result in the same toxicity as if administered alone. We believe that considering all combinations of dose levels between the two agents as a possible MTD should be acceptable, under medical restrictions and prior knowledge of such combinations. The recommended combination at the end of the trial should not be limited to the combination of both single-agent MTDs, but the dose-finding process should be performed similarly to that of a single-agent in order to recommend the combination with a toxicity rate closest to or below a pre-defined target. Indeed, in the same way, combining two agents can also induce a synergistic, antagonistic, or independent effect on overall efficacy. This point should be discussed for each combination of drugs, as the mechanism of action of each agent can differ.

In conclusion, we believe that the design of drug-combination phase I trials in oncology can be improved. We recommend that the starting doses of the drugs, as well as the dose levels and the dose-escalation steps, need to be appropriately justified. These parameters should be determined with the aim to: 1) ensure patient safety; 2) treat as few patients as possible at presumably infra-therapeutic doses; and 3) identify the optimal drug combination for further evaluation. We strongly support the use of innovative designs that are able, at least in theory, to fulfil these requirements.

Figure legends

Figure 1.

(a) Partial known ordering between combinations. (b) Possible orderings between combinations according to increasing DLT rates.

(a)



(b)

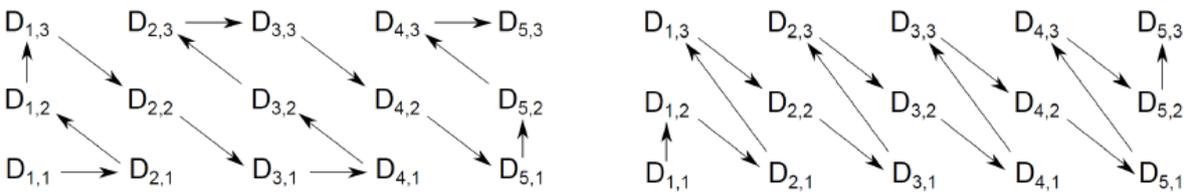
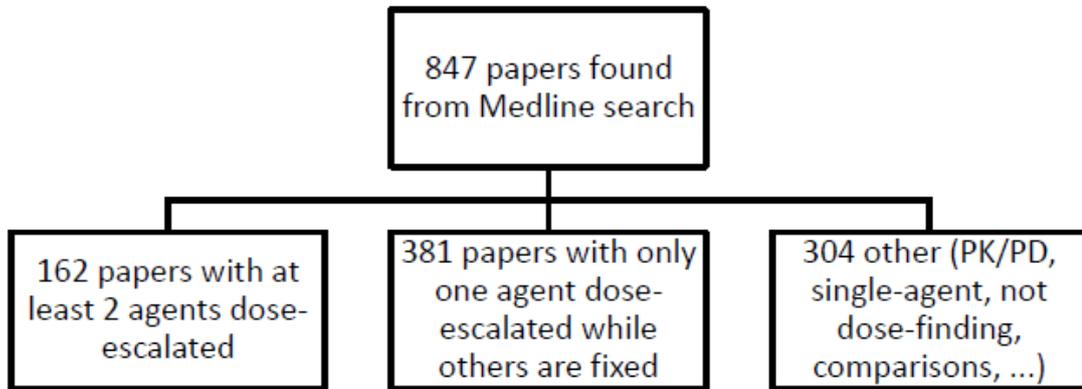


Figure 2.

Flow chart of the publications found from the MEDLINE® PubMed search.



PK=Pharmacokinetic, PD=Pharmacodynamic

Figure 3.

Control quality of 162 trials reviewed.

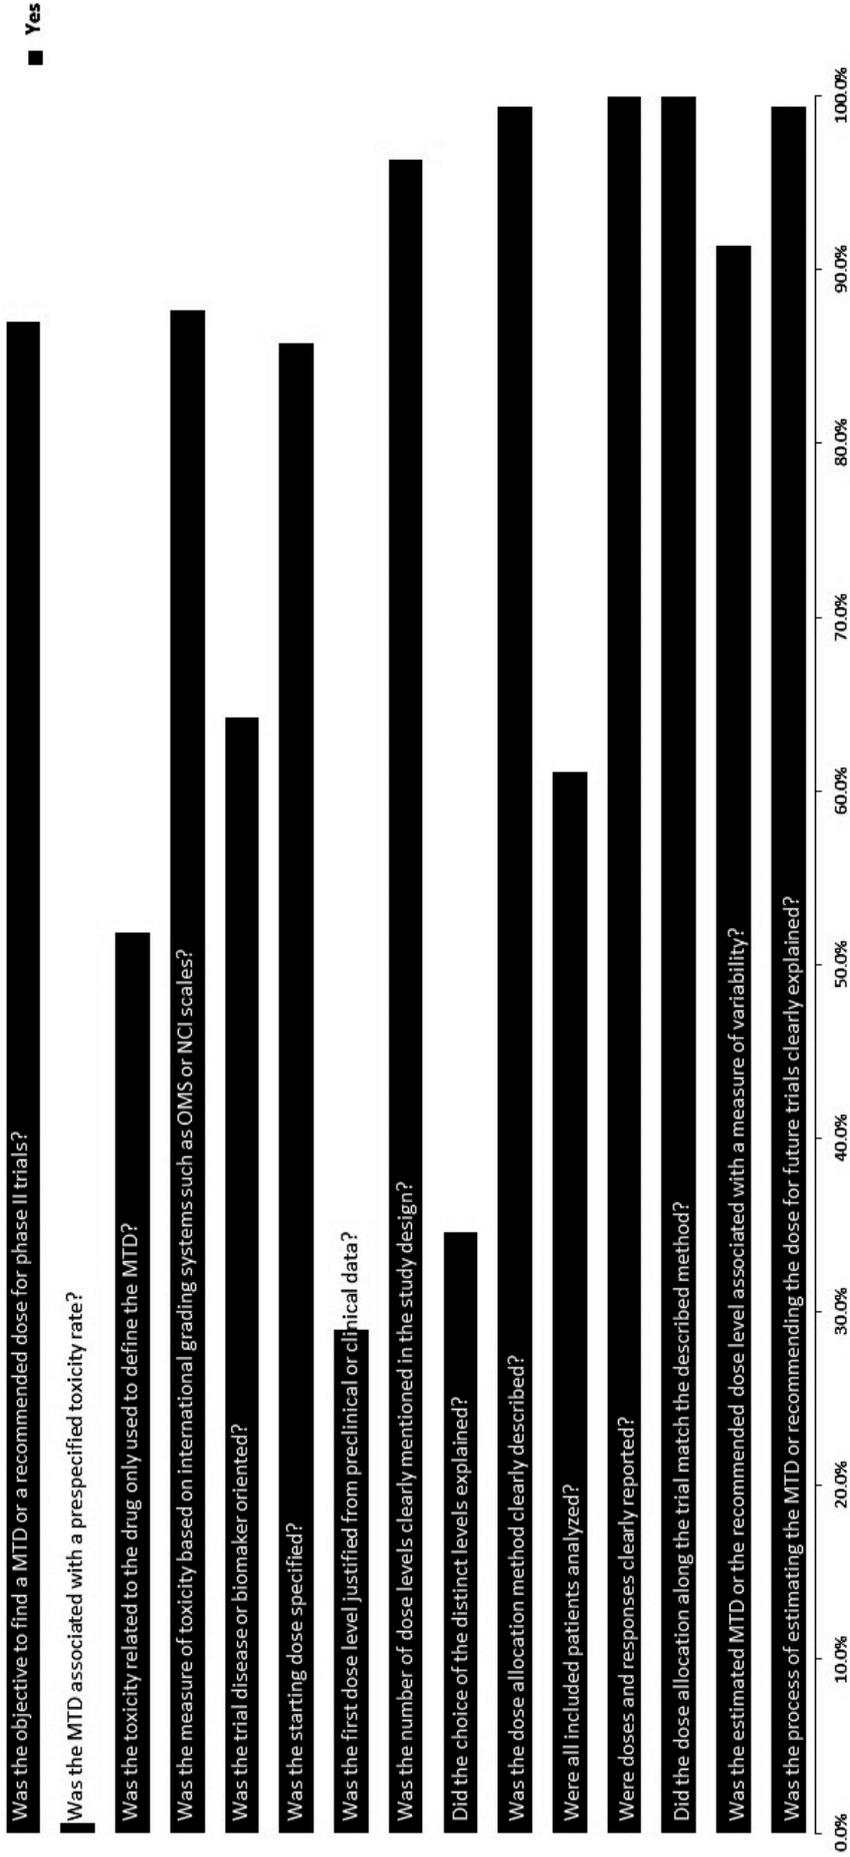


Table 1. Characteristics of the 162 drug combinations phase I trials reviewed

	<i>N</i> (%)	Median [range]
Number of drugs undergoing dose escalation:		
- 2	147 (91%)	
- 3	14 (9%)	
- 4	1 (1%)	
Types of drugs undergoing dose escalation:		
- cytotoxic agents only	55 (34%)	
- MTAs only	43 (27%)	
- combination of cytotoxic agent(s) and MTA(s)	64 (39%)	
Median number of patients per trial [range]		25 [7-136]
Starting doses of the drugs:		
- lowest combination of the trial	112 (69%)	
- Higher combination	50 (31%)	
Median number of dose combination levels considered:		
- 2-drug combinations		5 [2-16]
- 3-drug combinations		5.5 [3-15]
- 4-drug combinations		12 [12-12]
Ratio of the number of planned combinations to the number of possible combinations:		
- 2-drug combinations		0.67 [0.25-1]
- 3-drug combinations		0.24 [0.17-0.63]
- 4-drug combinations		0.13 [0.13-0.13]
Addition of intermediate dose levels during the trial:		
- Yes	38 (24%)	
- No	124 (77%)	
Dose-escalation design used:		
- 3+3 or modified 3+3 algorithm-based design	142 (88%)	
- model-based design	7 (4%)	
- combination design	1 (1%)	

MTA = Molecularly targeted agent

Acknowledgments of research support:

This work was partially funded by grants from the ANRT (Association Nationale de la Recherche et de la Technologie), “Laboratoires Servier”(CIFRE number 2011/0900).

Disclosure:

M-K. Riviere and F. Dubois are employees of IRIS (Institut de Recherches Internationales Servier) Pharmaceutical industry. All remaining authors have declared no conflicts of interest.

References

- [1] Storer BE. Design and analysis of phase I clinical trials. *Biometrics* 1989; 45: 925–37.
- [2] Cheung YK: Dose Finding by the Continual Reassessment Method. Chapman & Hall/CRC, 2011
- [3] Le Tourneau C, Lee JJ, Siu LL. Dose escalation methods in phase I cancer clinical trials. *J Natl Cancer Inst* 2009; 101: 708-20.
- [4] Curigliano G, Spitaleri G, Magni E, et al. Cisplatin, etoposide and continuous infusion bleomycin in patients with testicular germ cell tumors: efficacy and toxicity data from a retrospective study. *J Chemother* 2009; 21(6):687-92.
- [5] Bamias A, Aravantinos G, Kastriotis I, et al. Report of the long-term efficacy of two cycles of adjuvant bleomycin/etoposide/cisplatin in patients with stage I testicular nonseminomatous germ-cell tumors (NSGCT): a risk adapted protocol of the Hellenic Cooperative Oncology Group. *Urol Oncol* 2011; 29(2):189-93.
- [6] Slamon DJ, Leyland-Jones B, Shak S, et al. Use of chemotherapy plus a monoclonal antibody against HER2 for metastatic breast cancer that overexpresses HER2. *N Engl J Med* 2001; 344(11):783-92.
- [7] Flaherty KT, Infante JR, Daud A, et al. Combined BRAF and MEK inhibition in melanoma with BRAF V600 mutations. *N Engl J Med* 2012; 367(18):1694-703.
- [8] Soria JC, Massard C, Izzedine H. From theoretical synergy to clinical supra-additive toxicity. *J Clin Oncol*. 2009; 27:1359-61.
- [9] Mandrekar SJ. Dose-finding trial designs for combination therapies in oncology. *J Clin Oncol* 2014; 32:65-7.
- [10] Adjei AA. What is the right dose? The elusive optimal biologic dose in phase I clinical trials. *J Clin Oncol* 2006; 24(25):4054-5.
- [11] Cannistra SA. Challenges and pitfalls of combining targeted agents in phase I studies. *J Clin Oncol* 2008 ; 26(22):3665-7.
- [12] Hamberg P, Ratain MJ, Lesaffre E, et al. Dose-escalation models for combination phase I trials in oncology. *Eur J Cancer* 2010; 46(16):2870-8.

- [13] Harrington JA, Wheeler GM, Sweeting MJ, et al. Adaptive designs for dual-agent phase I dose-escalation studies. *Nat Rev Clin Oncol* 2013; 10(5):277-88.
- [14] Jain RK, Lee JJ, Hong D, et al. Phase I oncology studies: evidence that in the era of targeted therapies patients on lower doses do not fare worse. *Clin Cancer Res* 2010; 16(4):1289-97.
- [15] LoRusso PM, Boerner SA, Seymour L. An overview of the optimal planning, design, and conduct of phase I studies of new therapeutics. *Clin Cancer Res* 2010; 16(6):1710-8.
- [16] Booth CM, Calvert AH, Giaccone G, et al. 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). *Eur J Cancer* 2008; 44(1):19-24.
- [17] Seymour LK, Calvert AH, Lobbezoo MW, et al. Design and conduct of early clinical studies of two or more targeted anticancer therapies: recommendations from the task force on Methodology for the Development of Innovative Cancer Therapies. *Eur J Cancer* 2013; 49(8):1808-14.
- [18] U.S. Department of Health, Human Services, F., Drug Administration, C. f. B. E., and Research. Guidance for Industry, Clinical Considerations for Therapeutic Cancer Vaccines.
- [19] Le Tourneau C, Gan HK, Razak AR, et al. Efficiency of new dose escalation designs in dose-finding phase I trials of molecularly targeted agents. *PLoS ONE* 2012; 7(12):e51039.
- [20] Thall P F, Cook J D. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* 2004; 60 (3):684–693.
- [21] Doussau A., Thiebaut R., Paoletti X. Dose-finding design using mixed-effect proportional odds model for longitudinal graded toxicity data in phase I oncology clinical trials. *Stat Med* 2013; 32 (30):5430–5447.
- [22] Zhang W., Sargent D J, Mandrekar S. An adaptive dose-finding design incorporating both toxicity and efficacy. *Stat Med* 2006; 25:2365–2383.
- [23] Huang X, Biswas S, Oki Y, et al. A parallel phase I/II clinical trial design for combination therapies. *Biometrics* 2007; 63: 429-36.

- [24] Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics* 2005; 61: 217–22.
- [25] Mandrekar SJ, Cui Y, Sargent DJ. An adaptive phase I design for identifying a biologically optimal dose for dual agent drug combinations. *Stat Med* 2007; 26: 2317-30.
- [26] Mandrekar SJ, Qin R, Sargent DJ. Model-based phase I designs incorporating toxicity and efficacy for single and dual agent drug combinations: methods and challenges. *Stat Med* 2010; 29: 1077–83.
- [27] Yin G, Yuan Y. A latent contingency table approach to dose finding for combinations of two agents. *Biometrics* 2009; 65: 866–75.
- [28] Yin G, Yuan Y. Bayesian dose finding in oncology for drug combinations by copula regression. *Applied Statistics JRSS* 2009; 58: 211–24.
- [29] Thall PF, Millikan RE, Mueller P, et al. Dose-finding with two agents in Phase I oncology trials. *Biometrics* 2003; 59: 487–96.
- [30] Riviere MK, Dubois F, Zohar S. Competing designs for drug combination in Phase I dose-finding clinical trials. *Stat Med* 2014. [Epub ahead of print] doi: 10.1002/sim.6094.
- [31] Ivanova A, Wang K. A non-parametric approach to the design and analysis of two-dimensional dose-finding trials. *Stat Med* 2004; 23: 1861–70.
- [32] Gandhi L, Bahleda R, Tolaney SM, et al. Phase I study of neratinib in combination with temsirolimus in patients with human epidermal growth factor receptor 2-dependent and other solid tumors. *J Clin Oncol* 2014; 32: 68-75.
- [33] Conaway MR, Dunbar S, Peddada SD. Designs for single- or multiple-agent phase I trials. *Biometrics* 2004; 60:661-9.
- [34] Jones DR, Moskaluk CA, Gillenwater HH, et al. Phase I trial of induction histone deacetylase and proteasome inhibition followed by surgery in non-small-cell lung cancer. *J Thorac Oncol* 2012; 11:1683-90.

Chapter 2

State of the art: Dose-finding designs for single-agent trials in oncology

2.1 Standard designs for single-agent dealing with toxicity

Many designs for single-agent trials have been developed in the recent years [94]. They can be classified as algorithm-based approaches and model-based approaches [65].

2.1.1 Algorithm-based designs

Algorithm-based designs are non-parametric methods for which patients' allocation is based on a pre-defined algorithm according to the observed toxicity data [16]. They are also called “up-and-down” designs as dose levels can be either escalated or deescalated for the next cohort of patients according to decision rules based on the observed number of DLTs. The toxicity probabilities are often estimated empirically with observed toxicity rates. Most of the time, these algorithms are “short-memory” as they do not use all the information available but only the toxicity responses on the last cohorts [57]. Until recently, they were considered as “standard” designs and widely used in phase I clinical trials as they are easily understood by the medical community and easy to implement. Many authors have contributed to propose several algorithm-based designs [67, 39, 38]. In the following paragraphs, we will present three majors and well-known methods: the “3+3”, “A+B” and accelerated titration designs. Durham, Flournoy, and Rosenberger [20] described a family of random walk rules for the sequential allocation of dose levels.

2.1.1.1 Traditional “3+3” design

The most traditional design is algorithmic-based and called “3+3” [67] because patients are included by cohort of 3. The first cohort is treated at the lowest dose and then doses are escalated step by step. After treating 3 patients at a dose level, the next 3 patients are treated at the next higher dose level if no patient has encountered a DLT, while the same level is repeated if 1 DLT was observed. If at least 2 DLTs were observed on the 3 or 6 patients, then the dose-finding is stopped and the recommended dose level is defined as the next lower dose level. The algorithm can be summarised by Figure 2.1.

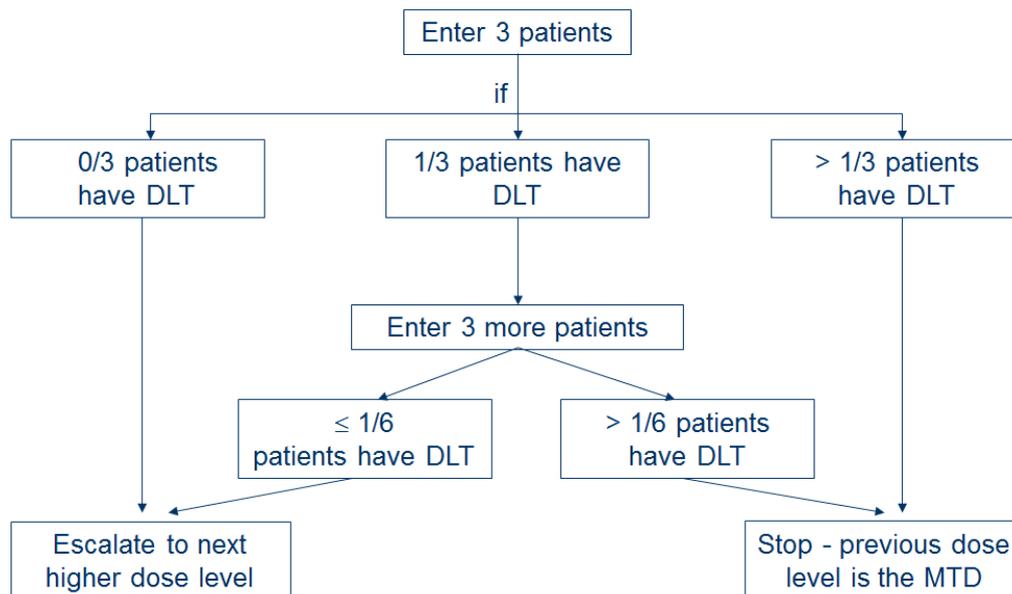


Figure 2.1: “3+3” strict traditional escalation rule.

One drawback of this design is that it is memory-less. The decision to escalate or stop the trial is only based on the last one or two observed cohort(s), ignoring everything happened in the previous cohorts [57]. The target (probability of) toxicity is not chosen but implicit, in this case it is inferior to 0.33. Moreover, you can never test again a dose which was already administered, thus you have at most 6 patients on a dose, which is very low. Therefore the exact confidence interval of Clopper-Pearson (see appendix A.2) at 95% for the probability of toxicity at a given dose is huge:

- [0; 0.71] if 0/3 toxicity was observed
- [0.09; 0.99] if 2/3 toxicities was observed
- [0.29; 1] if 3/3 toxicities were observed
- [0; 0.64] if 1/6 toxicity were observed

- [0.04; 0.78] if 2/6 toxicities were observed
- [0.11; 0.88] if 3/6 toxicities were observed
- [0.22; 0.96] if 4/6 toxicities were observed

In addition, with this design, when there is a mistake, the method does not take it into account. For instance after treated several patients or cohorts, if a patient was treated with a wrong dose level, or if a toxicity response is mis-evaluated, it would have changed all the dose allocation process with no possibility to take it into account in the dose recommendation.

Several simulation studies showed that this design tends to select dose under the true MTD [37, 64], especially when there are many doses and when the MTD is in the highest doses. Thus more patients are treated to low and possibly ineffective dose levels. Nevertheless, it is the most used design for many years and still may be, because this algorithm is easily understood by physicians and does not require the involvement of the statistician.

2.1.1.2 “A+B” scheme without dose de-escalation

The “A+B” up-and-down design proposed by Ivanova [38] is the generalization of the “3+3” algorithm. Indeed, the “3+3” is special case of “A+B” scheme with $A=B=3$ and $C=D=E=1$. The algorithm is summarized in Figure 2.2.

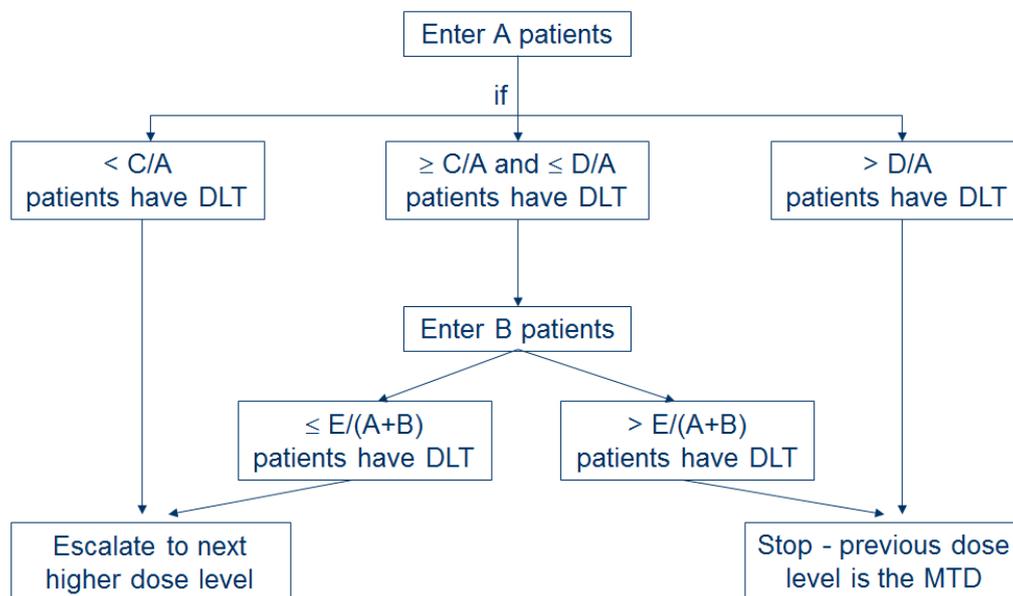


Figure 2.2: "A+B" escalation design.

This method is more flexible than the “3+3” as it enables to modulate the targeted toxicity. Indeed, depending of the type of cancer, a 33% toxicity target is

not always suitable. Moreover, the number of patients to include can be adjusted depending on the desired compromise between estimation reliability and the speed of the trial progress.

2.1.1.3 Accelerated titration designs

Simon et al. [66] have also introduced “accelerated titration designs” that are based on a rapid initial dose-escalation phase and on intra-patient dose escalation. The dose allocation process is based on DLTs and moderate toxicity. The most frequently used accelerated titration design can be described as follows: while no DLT or less than two moderate toxicities are experienced, patients are included by cohort of one and the dose level is escalated by steps of 40% increments (or 100%). Once the accelerated phase ends, a standard “3+3” dose-escalation scheme proceeds. This design enables to reduce the total number of patients that needs to be included in phase I and is one of the few methods with an intra-patient dose-escalation rule.

2.1.2 Model-based designs

Designs based on a model for cytotoxic phase I clinical trials have been developed in order to minimize the number of patients treated at unacceptably high toxic or at low inefficient dose levels. Model-based designs for single-agent trials have expanded since the publication of the continual reassessment method (CRM) by O’Quigley, Pepe and Fisher [54] in 1990. Several modifications of the CRM have been proposed [23, 30, 10]. Unlike algorithm-based methods, the toxicity probabilities are estimated using a statistical model using all the information accumulated along the trial. Many alternative dose-finding procedures have been proposed resulting in a growing statistical literature for model-based methods [81, 82, 5]. The literature describing model-based designs is very large and is well summarized in the reviews and collections of articles [65, 13, 75]. In the following sections, we will describe two important contributions and most used model-based designs: the continual reassessment method and the escalation with overdose control (EWOC).

2.1.2.1 Continual Reassessment Method

The CRM, introduced by O’Quigley et al. [54] as an alternative to the drawbacks of the “3+3”, has been designed to provide an ethical dose-allocation strategy that optimizes the proportion of patients treated at doses sufficiently close to the MTD while treating fewer patients with dose levels outside the therapeutic window. At each sequential analysis, all the accumulated DLT data are used for the estimation. Several modifications of the CRM were proposed, but the design was proposed

originally in a Bayesian framework with a one-parameter model. Garret-Mayer [26] published a detail review on the CRM.

Let d_k ($k = 1, \dots, K$) denote the dose levels investigated in a clinical trial. The dose administered to the i^{th} entered patient, X_i ($i = 1, \dots, n$) can be thought of as random variable, taking values $x_i \in \{d_1, \dots, d_K\}$. Let Y be a binary random variable, where 1 denotes the observation of a dose limiting toxicity and 0 not. A working model, w_k , representing the prior toxicity probabilities at each dose level d_k and α_i corresponds to the prior toxicity probability of the dose level administered to patient i . A monotonically increasing function $\pi(d, \beta)$ is used to model the dose-toxicity relationship $P(Y = 1|d)$, where β is a vector of parameters to estimate. A target toxicity is chosen (for example 0.3 on Figure 2.3). The number of patients per cohort is variable and must defined prior to the trial onset. Smaller cohort size provide better operational characteristics for the CRM.

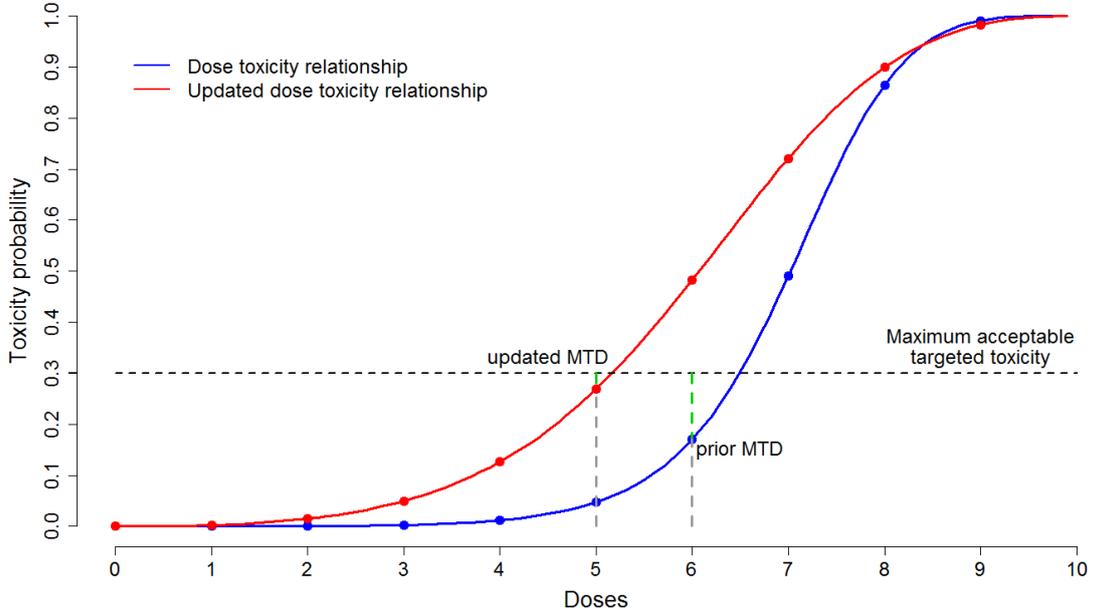


Figure 2.3: Continual Reassessment Method (CRM).

After each new patient or each new cohort of patients have completed the number of treatment cycles required, the likelihood is given by:

$$L(\beta|\text{data}) = \prod_{i=1}^I \pi(x_i, \beta)^{y_i} (1 - \pi(x_i, \beta))^{(1-y_i)}$$

enabling computation of the posterior density for the vector of parameters a given the joint prior density g of β :

$$f(\beta|\text{data}) = \frac{L(\beta|\text{data})g(\beta)}{\int_{\mathbb{A}} L(\beta|\text{data})g(\beta)d\beta}$$

Posterior toxicity probabilities, $\tilde{\pi}$, at each dose level are updated either (1) according to the toxicity model taken in the posterior mean of all parameters β , or (2) according to the posterior mean of the toxicity model. Then, the next cohort of patients is treated at the dose level corresponding to an estimated toxicity probability closest to the target toxicity:

$$\operatorname{argmin}_{1 \leq k \leq K} |\tilde{\pi}(d_k, \beta) - \theta|$$

For instance, in Figure 2.3, before cohort inclusion the prior MTD was dose level 6 as the corresponding toxicity probability minimizes the distance to the target toxicity. After cohort inclusion and re-estimation of the dose-toxicity relationship, the updated MTD is dose level 5.

The trial ends when the maximum fixed sample size is reached or if a pre-specified stopping rule for early termination is achieved. Cheung and Chappell [11] showed that under certain conditions (including large enough sample size), this method converges to the true MTD. Nevertheless, in phase I oncology, the number of patients is small (often 25 to 30 patients). However, even with a small sample size, CRM has better operating characteristics than the “3+3” [37, 64]: it decreases the number of patients treated under the MTD; a very reasonable number of patients are treated over the MTD; and the number of patients treated at the MTD is high. Stopping rules can be added to stop the trial earlier either if enough information are collected to declare having identifying the MTD, or to stop the trial for over- or under-toxicity. Moreover, in opposite to the “3+3”, the CRM enables to take into account several types of mistakes as the correction of the dose level truly administered to patients or DLTs re-evaluation.

Withal, Neuenschwander et al. [53] raised some concerns about the CRM. Indeed, they have highlighted some reasons explaining the performance failure of the CRM in some cases: (1) the use point estimates instead of more informative posterior summaries, (2) the regular use of a less flexible one-parameter instead of a two-parameter model, (3) the choice of working model, and (4) the coupling of the Bayesian procedure with a mis-specified model and/or suboptimal dose selection rules.

Choice of the working model

O’Quigley and Zohar [58] studied the choice of the working model and underlined (among others) that an adequate spacing between prior toxicity probabilities is important. In the same way, Lee and Cheung [47, 48] studied how to calibrate the working model in the context of phase I single-agent evaluation. They developed an approach in order to maximize the percentage of correct selection at the end of the trial when the physicians are not able to give the prior toxicity probabilities at each dose levels prior to the trial onset. They developed the “dferm” R package containing the “getprior” that enable the elicitation of the working model

given the prior position of the MTD, the number of dose levels, the target toxicity and the length of indifference intervals. For more detail about the calibration method proposed by Lee and Cheung [47], see Appendix A.1. Zhang and al. [90], Yuan and Yin [88], Daimon et al. [17] among others also tried to highlight how to calibrate the working model for the CRM.

1-parameter vs 2-parameter models

The use of one-parameter vs two-parameter models for the CRM was largely studied [60, 53]. Nevertheless, no consensus was found and the statistical community is still divided on the number of parameters to use to model the dose-toxicity relationship.

Several models were proposed and compared in the CRM, we present some examples below:

- 1-parameter (a) tangent model:

$$\pi(d_k, a) = \left(\frac{\tan(w_k) + 1}{2} \right)^a$$

with $a > 0$.

- 1-parameter (a) empirical/power model:

$$\pi(d_k, a) = w_k^a$$

with $a > 0$.

or

$$\pi(d_k, a) = w_k^{\exp(a)}$$

with $a \in \mathbb{R}$.

- 1-parameter (a) logistic model ($b=3$):

$$\pi(d_k, a) = \frac{\exp(3 + a \times f(d_k))}{1 + \exp(3 + a \times f(d_k))}$$

with $a > 0$ so that the toxicity probability is increasing with the dose, and

$$f(d_k) = \log \left(\frac{w_k}{1 - w_k} \right) - 3.$$

or

$$\pi(d_k, a) = \frac{\exp(3 + \exp(a) \times f(d_k))}{1 + \exp(3 + \exp(a) \times f(d_k))}$$

with $a \in \mathbb{R}$ so that the toxicity probability is increasing with the dose, and

$$f(d_k) = \log \left(\frac{w_k}{1 - w_k} \right) - 3.$$

- 2-parameter (a, b) logistic model:

$$\pi(d_k, a, b) = \frac{\exp(b + a \times f(d_k))}{1 + \exp(b + a \times f(d_k))}$$

with $b \in \mathbb{R}$, and $a > 0$ so that the toxicity probability is increasing with the dose.

Several simulation studies were performed to give some guidance about the choice of one-parameter or two-parameter models giving contradictory recommendations as results depend on the chosen scenarios. Some statisticians recommend the use of one-parameter model as they judge that due to the small sample size, the estimation of two parameters is not reliable. Paoletti and Kramar [60] also explore the use of these two types of model with both likelihood and Bayesian approaches and showed that, on average, the performances of a one-parameter model is superior and that the power model has good operating characteristics. Nevertheless, as stated previously, Neuenschwander [53] highlighted drawbacks of the CRM in the case of one-parameter models and recommend the utilization of two-parameter models that Novartis uses in its current practice. Chevret [12] studies extensively the impact of the functional form of the model and the prior distribution of its parameter on the performance of the method.

Bayesian vs Frequentist

Six years after the first paper on the CRM, O’Quigley and Shen [56] introduced a non-Bayesian version of the CRM: the likelihood CRM also called CRML. Both version of the CRM, Bayesian or Frequentist, give in general similar results. Paoletti and Kramar [60] carried out a large simulation study where they compared the CRM with both a likelihood approach and a Bayesian approach for model estimations.

Comparison “3+3” vs CRM

Several authors have compared in details the performance of the “3+3” and the CRM, especially Iasonos et al . and Rogatko et al. [37, 64]. As mentioned above, even if several points are still debated for the CRM and model-based designs, the statistical community agrees that the CRM has better operating characteristics than the “3+3” and “3+3” should not be used to perform phase I clinical trials.

As a small example, we performed simulations on five different scenarios, two with five dose levels and three with ten dose levels. As it is usually done, the sample size for the CRM was chosen as $6 \times k$ (with k the number of dose levels in the trial) as six correspond to the number of patients included at a dose in the “3+3” when observing one DLT in the first cohort. It should be noted that for this

comparison, more patients may be enrolled with the CRM, but this design should be used with a reasonable number of patients, and the sample size of the “3+3” cannot be determined in advance.

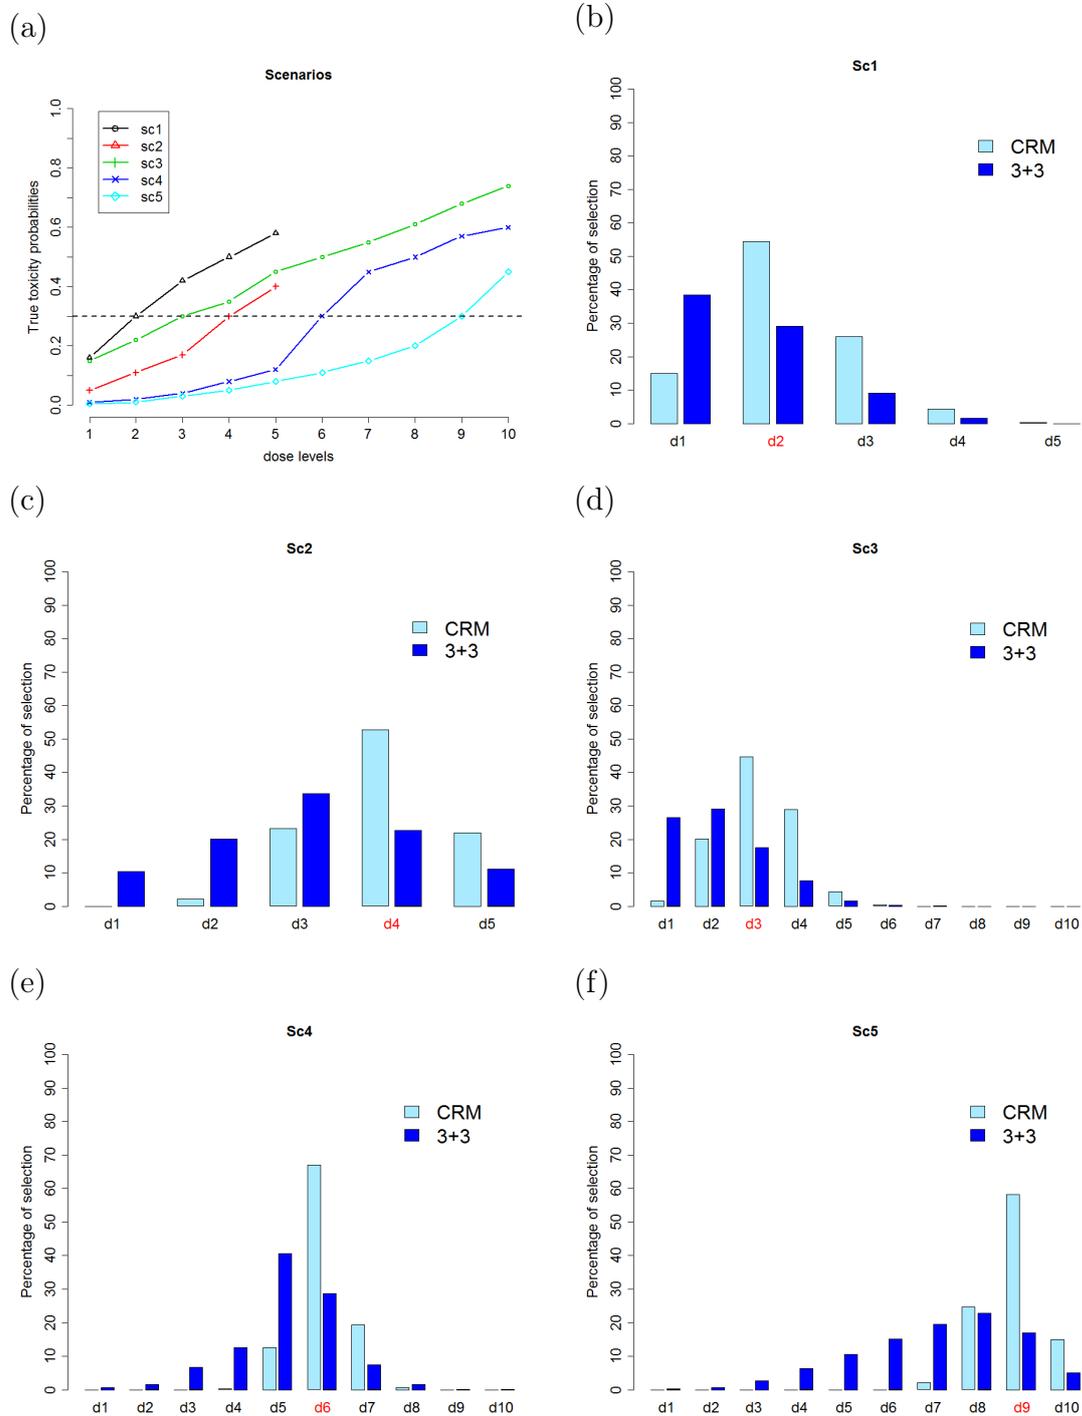


Figure 2.4: (a) Scenarios and (b-f) Comparison of percentage of selections for the “3+3” and the CRM. The true MTD is given in red.

We can observe that the “3+3” tends to under-estimate the MTD, especially by selecting the dose before the true MTD. The distribution of percentage of selection is spread. Percentage of correct selections for the “3+3” are low when there are many dose levels and the MTD is in the highest dose levels, as this method has difficulties to escalate. This is due to the fact that the method stop when 2 DLTs over 3 or 6 patients has been observed at one dose level. Indeed, for instance for scenario 5, given the true toxicity probabilities, we can calculate explicitly the probability to stop the trial before achieving the true MTD, dose level 9. This probability is equal to 0.54. Therefore, due to the small number of patients, and due to the repetition of the allocation rule at each dose level, the probability to observed unexpected DLTs increased, and the finally in this scenario, we have more than 1 risk on 2 to stop the trial before achieving the true MTD.

Other criteria than percentage of correct selection should also be considered in the comparison, such as the number of patients enrolled at the MTD, the number of observed toxicities, etc...

Stopping rules

In order to optimize the sample size and then to better address ethical concerns, several authors developed efficient stopping rules that allow to continue patient accrual or to stop inclusions. Zohar and Chevret [93] proposed Bayesian stopping rules for the CRM, based on either posterior or predictive probability distributions. These rules aim at early termination of the trial either for the mis-specification of the dose range leading to over-toxicity, or for finding the MTD according to a prefixed gain in the point estimate or accuracy in the estimated toxicity probability at the MTD. As an example, from [93], Zohar recommends the following stopping rule for practical use. This rule considers two actions, that is “continue the trial” and “stop the inclusions for non-usefulness of continuation” because a predetermined suitable estimation of the MTD has been reached. The idea of this rule is that: if the dose estimated to be the MTD has no expected gain by trial continuation (that is by including z additional patients in the trial) in estimate precision of associated toxic probability, then the trial should be stopped. According to [93], z can be chosen equal to 3. For these z new patients, at each individual patient inclusion the dose assumed to be administered is the MTD determined by the model. This stopping rule can be detailed as follows, if d^{MTD} is the dose estimated to be the MTD at that point in the trial,

$$\begin{cases} \text{Stop the trial if } G(z, d^{\text{MTD}}) < \xi \text{ occurs } t \text{ times} \\ \text{Continue otherwise} \end{cases}$$

where ξ is a pre-specified threshold to calibrate, and $G(z, d^{\text{MTD}})$ is the maximal predictive gain of z additional patients inclusions on the credibility interval width, c_α , of the toxicity probability of the recommended dose d^{MTD} defined as:

$$G(z, d^{\text{MTD}}) = \max_{(y_1, \dots, y_z)} | c_{\alpha, j+z}(f_{\pi, d^{\text{MTD}}}) - c_{\alpha, j}(f_{\pi, d^{\text{MTD}}}) |$$

with $f_{\pi, d^{\text{MTD}}}$ the posterior distribution for the toxicity probability π at dose d^{MTD} .

O'Quigley and Reiner [55] have also proposed a simple and efficient stopping rule whose idea is that continuing the study (that is including z additional patients in the trial) would not lead to a change in dose recommendation with high probability. The probability not to have any change in dose recommendation by inclusion of z additional patients is calculated from a probability tree (see Figure 2.5),

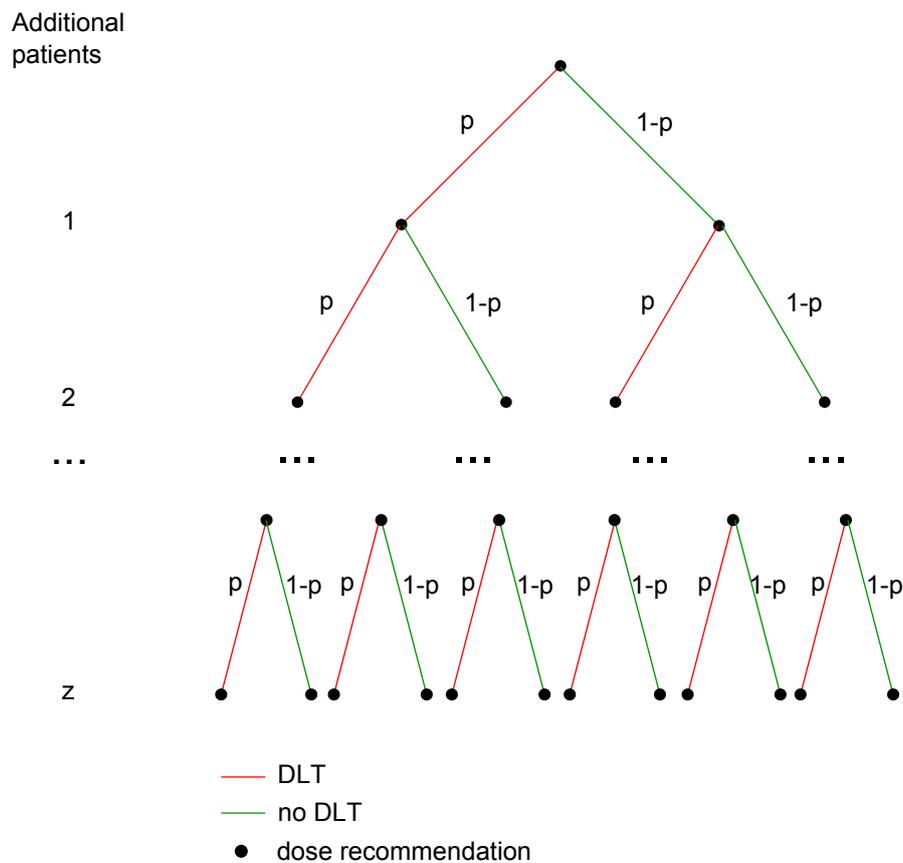


Figure 2.5: Probability tree for dose recommendation of z additional patients.

and can be detailed as follows:

$$\sum_{r=0}^k A_r (1-p)^{z-r} p^r$$

where A_r is the number of combinations (paths or branches of the tree) with r DLTs and without change in dose level recommendation along the z additional patients (from the dose level estimated to be the MTD after the inclusion of the

real data in the trial, i.e. before the inclusion of the fictitious z additional patients), and p is the estimated probability of having a DLT at that dose level (before the inclusion of the fictitious z additional patients). This probability of no change in dose recommendation is compared with a threshold ξ (for instance 0.90); if superior, the trial will end earlier for finding the MTD.

A practical example

Even if the “3+3” remains the most implemented design in past years, several phase I trials were set up using model-based designs, especially the CRM. In this paragraph, we present the results of a real phase I clinical trial using the CRM published in the British Journal of Cancer [49] (see Table 2.1). We are detailing step by step the statistical dose-finding process in the context of a real study. The aim of the trial was to assess the MTD of semisynthetic homoharringtonine (ssHHT) in the treatment of patients with advanced myeloid leukaemia. This trial involved 5 dose levels, d_k , among which dose level 4 was assumed to be the prior MTD for a target toxicity of 0.33. Initial guessed (prior) toxicity probabilities, w_k , were initiated to (0.05, 0.10, 0.15, 0.33, 0.50) prior to the trial onset according to investigators’ experience and to literature. The design allowed dose skipping to assign patients to the dose level associated with a toxicity probability closest to the target. The dose-toxicity model was chosen as a 1-parameter logistic model, with the intercept fixed to 3:

$$\pi(d_k, a) = \frac{\exp(3 + a \times f(d_k))}{1 + \exp(3 + a \times f(d_k))}$$

where $f(d_k) = \log\left(\frac{w_k}{1 - w_k}\right) - 3$ inducing the following dose-transformation from the prior toxicity probabilities, $f(d) = (-5.94, -5.20, -4.73, -3.71, -3.00)$, and $a > 0$. An exponential prior distribution with parameter 1 was chosen for the unknown parameter a , i.e., $a \sim \text{Exp}(1)$. Patients were included in cohorts of size 3. The trial could stop either if the fixed sample size was reached, or when stopping rules measuring futility of trial continuation were fulfilled [93].

The first cohort of three patients received the lowest dose level. No DLT was observed, and then posterior toxicity probabilities, $\tilde{\pi}$, were updated by solving the model in the estimation of the posterior value of parameter a , that is $\pi(d, \tilde{a})$. According to the model, the dose level estimated to be the closest to the target was dose level 5 with a DLT probability of 0.11. Nevertheless, for ethical reasons, investigators preferred not to skip up to the highest dose level 5, but to dose level 3. Then the dose allocation process continue by assigning patients to the dose level closest to the target toxicity. After six cohorts were included, the four last cohorts received the same dose level 4, and the decision was made by the expert committee to stop the trial with three over four stopping criteria detecting futility of trial continuation.

Table 2.1: Updated toxicity probabilities of the five dose levels after each newly included cohort using the CRM in the context of a real clinical trial.

				Dose levels				
				1	2	3	4	5
				Prior toxicity probabilities				
				0.05	0.10	0.15	0.33	0.50
Cohort	Dose	DLT/Patients	\tilde{a}	Posterior toxicity probabilities				
1	1	0/3	1.70	0.001	0.003	<u>0.006</u>	0.035	0.11
2	3	1/3	0.93	0.07	0.14	0.19	<u>0.39</u>	0.55
3	4	1/3	0.94	0.07	0.13	0.19	<u>0.38</u>	0.54
4	4	0/3	1.08	0.03	0.07	0.11	<u>0.27</u>	0.45
5	4	1/3	1.05	0.04	0.08	0.12	<u>0.29</u>	0.46
6	4	2/3	0.96	0.06	0.12	0.17	<u>0.36</u>	0.53

Time-to-event continual reassessment method

Several modifications and extensions of the CRM were proposed in adaptive dose-finding studies. A famous and practical extension is the time-to-event continual reassessment method (TITE-CRM) introduced by Cheung and Chappell [10] in order to deal with late-onset toxicities. Indeed, in practice a longer follow-up time can be needed to assess the DLT outcome. Therefore, the toxicity outcomes of some patients already treated in the trial may be unobserved (or censored) when the dose level should be re-evaluated to be assigned to a new cohort of patients in the trial. Waiting to assess DLTs for each entire cohort before including a new one in the trial can highly increase the duration of the trial. Thus, Cheung and Chappell [10] have developed the TITE-CRM to overcome this issue. Let T be a (maximum) fixed time window during which patients are followed, and $y_{i,N}$, $C_{i,N}$, and $\hat{w}_{i,N}$, respectively, be the indicator of a DLT for patient i prior to the entry of the $(N + 1)^{\text{th}}$ patient, the follow-up time of patient i prior to the entry of the $(N + 1)^{\text{th}}$ patient, and the weight assigned to patient i prior to the entry of the $(N + 1)^{\text{th}}$ patient. Let t_i denote the time to toxicity of the i^{th} patient. They proposed to consider a weighted dose-toxicity relationship $\hat{w}\pi(d, \beta)$, where \hat{w} are weights that are monotone and increasing with patients follow-up time and such that $0 \leq \hat{w} \leq 1$. Depending on the expected distribution of DLTs appearance, different weights should be chosen. For instance, Cheung and Thall [12] proposed weights $\hat{w}_{i,N}$ such as:

$$\hat{w}_{i,N} = \frac{\#\{m/X_m \leq C_{i,N}, C_{m,N} \geq T\} + \hat{w}_{i,N}^0}{\#\{m/X_m \leq T, C_{m,N} \geq T\} + 1},$$

where m refers to patient, $\#\{m/X_m \leq T, C_{m,N} \geq T\}$ is the number of completely followed patients on who a toxicity was observed, and $\hat{w}_{i,N}^0$ is the linear weight for

patient i defined as $\hat{w}_{i,N}^0 = C_{i,N}/T$, that is the proportion of time patient i was followed compared to the full follow-up time. The likelihood for the TITE-CRM with the weighted dose-toxicity model becomes:

$$L(\beta|\text{data}) = \prod_{i=1}^N (\hat{w}_{i,N}\pi(x_i, \beta))^{y_{i,N}} (1 - \hat{w}_{i,N}\pi(x_i, \beta))^{1-y_{i,N}}$$

Simulation studies have showed that this design performs well for late toxicities. Then, among others authors, Polley [61] proposed practical modifications to the TITE-CRM with fast patient accrual and late-onset toxicities.

Longitudinal data

Although in statistical analysis toxicity is considered as a binary outcome only on the first cycle of treatment, in practice toxicity is repeatedly measured over cycles on an ordinal scale using toxicity grades. Recently Doussau, Thiébaud and Paoletti [18] have proposed an adaptive dose-finding design using longitudinal measurements of ordinal AEs with a proportional odds mixed-effect models. The optimal dose is then the dose producing a target toxicity rate per cycle. Their model can also be used to identify cumulative or late toxicities.

2.1.2.2 Escalation With Overdose Control

In 1998, Babb, Rogatko and Zacks [5] developed in the context of cancer phase I clinical trials a design based on dose escalation with overdose control (EWOC). The probability of DLT at a given dose level, $P(Y = 1|d)$, is given by the model

$$\pi(d, \beta = (\beta_0, \beta_1)) = F(\beta_0 + \beta_1 d),$$

where F is a specified link function and $\beta_1 > 0$ so that the toxicity probability is increasing with the dose. The MTD, d^{MTD} , is the dose level such that the toxicity probability is equal to θ . Therefore

$$d^{\text{MTD}} = \frac{F^{-1}(\theta) - \beta_0}{\beta_1} = d_1 + \frac{F^{-1}(\theta) - F^{-1}(\rho_0)}{\beta_1},$$

where ρ_0 denotes the toxicity probability at the starting dose d_1 . In the same manner as the CRM, the likelihood and posterior joint distribution of the parameters can be calculated. However, the authors prefer to consider a re-parametrization of the model: $T(\beta_0, \beta_1) = (\rho_0, d^{\text{MTD}})$. Therefore the explicit joint posterior distribution, $\tilde{\pi}(\rho_0, d^{\text{MTD}}|\text{data})$ can be calculated (for formula see [5]) as well as the marginal posterior distribution:

$$\tilde{\pi}(d^{\text{MTD}}|\text{data}) = \int_{\rho_0} \tilde{\pi}(\rho_0, d^{\text{MTD}}|\text{data})d\rho_0,$$

Then the cumulative distribution function of the MTD in z is given by:

$$P(d^{\text{MTD}} \leq z | \text{data}) = \int_{d_1}^z \tilde{\pi}(d^{\text{MTD}} | \text{data}) \mathbf{1}_{d^{\text{MTD}} \in [d_1; d_K]} dd^{\text{MTD}},$$

The EWOC design consists in assigning the next patient or cohort of patients to the dose level x_k such that

$$P(d^{\text{MTD}} \leq x_k | \text{data}) = \alpha,$$

That is, the aim is to administer to patients the dose level corresponding to a posterior probability it exceeds the MTD equal to α . This allocation rule selects any continuous dose level between d_1 and d_K . Nevertheless, if due to logistic practical constraints often encountered, the dose range is a discrete pre-specified set of doses, the next dose level to assign is defined as follows:

$$\max_{1 \leq k \leq K} d_k, \{d_k - x \leq T_1 \text{ with } P(d^{\text{MTD}} \leq x | \text{data}) - \alpha \leq T_2\},$$

where $T_1, T_2 > 0$ are called tolerances.

At the end of the trial, the MTD is estimated by minimizing the posterior expected loss with respect to some loss function l , that is:

$$d^{\text{MTD}} = \underset{\gamma \in [d_1; d_K]}{\operatorname{argmin}} \int_{d_1}^{d_K} l(\gamma, d) \tilde{\pi}(d | \text{data}) dd.$$

Different loss functions should be considered asymmetric or not depending on the desired compromise between underestimation and overestimation.

Several modifications and extensions were also proposed for this design, as well as a software [74, 73, 9, 52]. Haines, Perevozskaya, and Rosenberger [31] proposed a c- and D-optimal Bayesian phase I design.

2.2 Designs for single-agent trials dealing with both toxicity and efficacy

Novel therapies such as MTA are challenging as they question the the standard settings of drug development. The assumptions for phase I trials in oncology for cytostatic agents are different from that of cytotoxic agents. For cytostatic agents, efficacy does not necessarily increase monotonically with the dose levels, but likely plateau after they reach maximal efficacy.

Several authors have proposed methods for dose-finding based on both efficacy and toxicity. Among other authors, we will briefly present some contributions. Thall and Russell [72] proposed a Bayesian adaptive design for conducting single-arm clinical trials for trinomial outcomes. They considered the following three outcomes: (1) no response (= no efficacy and no severe toxicity), (2) success (=efficacy and no severe toxicity), and (3) toxicity (= DLT). Thall and Russell used

a proportional odds model to fit both dose-toxicity and dose-efficacy relationships. Then, due to issues encountered in the practical use of this design, Thall and Cook [69] developed a new design. They proposed to partition the two-dimensional toxicity–efficacy probability domain by introducing a trade-off contour. The set of contours was constructed with a polynomial model based on three points derived from physicians’ knowledge. The next cohort of patients is treated at the best dose level according to all current accumulated data in the trial. Bekele and Shen [7] investigated a joint distribution of a binary and a continuous outcome by introducing latent variables in a probit model. Zhang et al. [91] proposed a modification of the design of Thall and Russel in an approach based on a flexible continuation-ratio model with different optimal dose selection criteria and different stopping rules. Yin et al. [84] developed a dose-finding Bayesian adaptive design for phase I/II clinical trials in order to incorporate both toxicity and efficacy. They jointly model the bivariate binary data to take into account the correlation between toxicity and efficacy. The dose allocation process is then based on an odds ratio criteria constructed from the posterior toxicity and efficacy probabilities. Finally Mandrekar et al. [51] reviewed two model-based designs utilizing a proportional odds model or a continuation ratio model with both independent toxicity and efficacy curves. They used a CRM with a simple dose selection criterion.

Chapter 3

State of the art for combination designs: Competing designs for drug combination in phase I dose-finding clinical trials

As combination therapies have become current practice in the recent years, new statistical methods are required for phase I clinical trials in oncology that take into account the multidimensionality of the problem. That is why after the 2000's, several authors have developed designs specific for combination studies. Thall et al. [70] proposed a Bayesian dose-finding method based on a six-parameter model. Ivanova and Wang [40] provided a up-and-down dose-allocation algorithm where the next combination is determined in the neighborhood according to the proportion of observed DLTs. A set of MTDs is selected after performing isotonic regression on the empirical proportions of DLTs. Conaway et al. [15] proposed a dose-finding method based on the simple and partial orders between toxicity probabilities of drug combinations. Wang and Ivanova [79] developed a three-parameter model-based method in which the parameters are estimated using Bayesian inference. Huang et al. [35] proposed a phase I/II design based on the "3+3" type dose escalation scheme. Mandrekar et al. [50, 51] proposed an approach incorporating the toxicity and efficacy of each agent into the identification of an optimal dosing region for the combination by using a continuation ratio model to separate each agent's toxicity and efficacy curves. Yuan and Yin [87] proposed a sequential dose-finding design that allows single-agent dose-finding methods to be used in multiple-agent combination trials. Fan et al. [22] proposed a two-stage "2+1+3" algorithm-based design. Yin and Yuan [85] developed a Bayesian adaptive design based on latent 2×2 tables in which the combination's toxicity probabilities in the two-dimensional space are estimated using a Gumbel-type model. Yin and Yuan [86] extended their method by changing to a copula-type model to simulate the effect of two or more drugs in combination. Bailey et al. [6] introduced a second agent as a covariate in a logistic model. Braun and Wang [8] proposed

a hierarchical-model-based approach for dose finding. Wages et al. [77, 78] considered an approach based on the continual reassessment method and taking into account different orderings with partial order between combinations. In this case, the MTD is estimated for the order associated with the highest model-selection criterion. Whitehead et al. [83] used a Bayesian approach based on an assumption of monotonicity in the relationship between the strength of the dose-combination and the distribution of the bivariate outcome.

Many designs have been proposed for dose-finding in the context of combination studies based on various methodological ideas. These designs have attractive theoretical concepts and simulation studies often showed high performance. Nevertheless, there was no review or recommendation concerning these methods and only very limited comparisons between two designs were available. A first part of my thesis consists of performing a simulation comparison among a state of the art of combination methods. We study each method through their models, start-up phases, estimations, allocation rules, ... Then, we select several combination designs representative of the statistical literature, and we slightly modify them to be comparable in a large simulation study. We computed ourselves each method for this comparison analysis. The detailed designs can be found in the paper attached below. The aim of our work is to provide statisticians involved in dose-finding studies with tools to evaluate combinations in order to select the most suitable design according to a clinical trial's combination and indication.

We found that in general, model-based designs seemed to perform better than algorithm-based designs, and between them they seemed comparable and no model-based design stood out from the others in our comparison. When selecting an ordered subset of combinations and applying the one dimensional CRM, the CRM seemed to have high performance when the true MTD is included in the combination pathway. (It should be noted that the same number of patients as combination designs was used for less combinations involved in the trial). Nevertheless, pre-selecting an arbitrary reduced subset of combinations induces a risk in selecting a combination with a DLT rate far from target toxicity. We do not know how this issue can be anticipated prior to the trial. We think that this point should be discussed for each combination of drugs, as the mechanism of action of each agent can differ, and that relies on physicians' knowledge. However, from our experience, it seems very difficult for physicians to determine the set of dose levels to be explored even in phase I single-agent trials. Therefore, determining the dose levels and also a subset of combinations to explore in combination trials seems even more complex. The success of a trial can depend on the statistical design, but also on many other parameters and at the root on the adequacy of the dose levels retained for the trial. Thus, if there is any doubt concerning the choice of combinations, we believe that all combinations should be considered and the design should allow the possibility of exploring any combination estimated to be the best.

This comparison study led to the writing of a manuscript published in the journal "Statistics in Medicine" that was very controversial. For comparison purposes, modifications of the designs were implemented. Moreover, the model-based

designs require the choice and the calibration of several parameters prior to the trial onset. The calibration of these parameters is subject to discussion and highly depend on the scenarios retained for the simulation studied. Despite the fact that we did not favor any method, and that they were all compared with the same choices and calibrations, this manuscript has just received two commentaries. It seems that convincing physicians to use combination designs is a big challenge due to the complexity and the ignorance of these methods, but it is also a challenge for statisticians as, even more than in single-agent trials, several settings are subject to discussions and debates.

Competing designs for drug combination in phase I dose-finding clinical trials

M-K. Riviere,^{a,b,*†} F. Dubois^b and S. Zohar^a

The aim of phase I combination dose-finding studies in oncology is to estimate one or several maximum tolerated doses (MTDs) from a set of available dose levels of two or more agents. Combining several agents can indeed increase the overall anti-tumor action but at the same time also increase the toxicity. It is, however, unreasonable to assume the same dose-toxicity relationship for the combination as for the simple addition of each single agent because of a potential antagonist or synergistic effect. Therefore, using single-agent dose-finding methods for combination therapies is not appropriate.

In recent years, several authors have proposed novel dose-finding designs for combination studies, which use either algorithm-based or model-based methods. The aim of our work was to compare, via a simulation study, six dose-finding methods for combinations proposed in recent years. We chose eight scenarios that differ in terms of the number and location of the true MTD(s) in the combination space. We then compared the performance of each design in terms of correct combination selection, patient allocation, and mean number of observed toxicities during the trials. Our results showed that the model-based methods performed better than the algorithm-based ones. However, none of the compared model-based designs gave consistently better results than the others. Copyright © 2014 John Wiley & Sons, Ltd.

Keywords: dose-finding studies; phase I trial; drug combination; Bayesian inference; oncology; cytotoxic

1. Introduction

The majority of phase I cytotoxic dose-finding studies seek to establish a dose high enough to be able to observe potential efficacy while maintaining the toxicity rate within certain pre-defined acceptable limits. In oncology, phase I studies focus on determining the maximum tolerated dose (MTD) that will be used in further phase II clinical trials of which the central interest is on potential efficacy. Following Storer [1], the statistical formulation of the problem is to select a dose level from several available doses, with a toxicity probability closest to a given target [2–4]. Most statistical model-based or algorithm-based methods were developed for cytotoxic single-agent phase I dose-finding clinical trials [5]. In this context, it is assumed that the toxicity of a single agent is monotonic and increases with the dose, as does the efficacy.

In the field of oncology, it is currently rare to find new molecules that perform better than existing therapeutic strategies. When combining two or more agents, there may be a potential synergistic effect in terms of efficacy. That is why investigators wish to increase overall anti-tumor action and survival by combining several agents, either cytotoxics or targeted molecules, or both. As a result, it is difficult to suppose that each molecule will act independently in terms of toxicity. In dose-finding studies, physicians aim to gradually increase toxicity during the dose-escalation procedure. However, when combining several agents, the ordering of toxicity probabilities is not fully known. For instance, the combination of two cytotoxics can induce an ordered subset of toxicity (Figure 1(a)). Even when a partial ordering is known, it is still difficult to decide how to escalate or de-escalate a combination of doses. Indeed, on a diagonal, there is no knowledge about which combination is more toxic; it is not known prior to the trial

^aINSERM, U1138, Equipe 22, Centre de Recherche des Cordeliers, Université Paris 5, Université Paris 6, Paris, France

^bIRIS (Institut de Recherches Internationales Servier), Suresnes, France

*Correspondence to: M-K. Riviere, INSERM, U1138, Equipe 22, Centre de Recherche des Cordeliers, Université Paris 5, Université Paris 6, Paris, France.

†E-mail: marie-karelle.riviere@crc.jussieu.fr

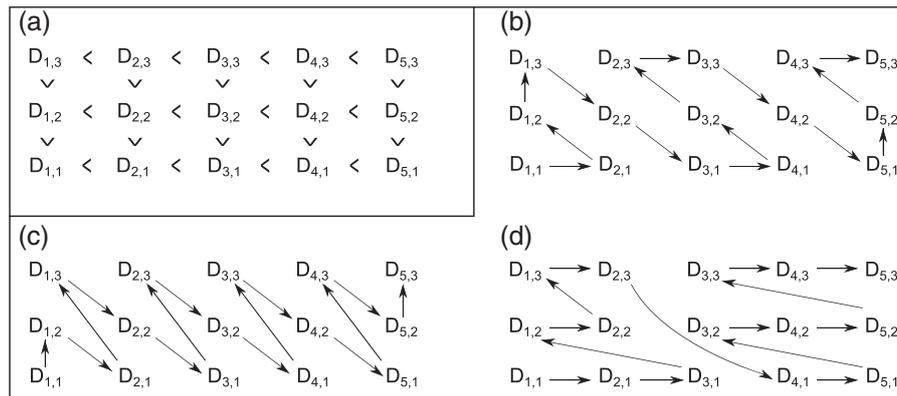


Figure 1. (a) Partial order relationships between combinations. (b) Three possible orderings consistent with the partial order retained for POCRM among all possible orderings for the simulation study.

which of $D_{1,2}$ and $D_{2,1}$ is more toxic. Therefore, it is senseless to use single-agent dose-finding methods in combination studies.

Several authors have recently addressed this issue by proposing new methods for combination studies, which are either algorithm-based or model-based. Ivanova and Wang have proposed an ‘up-and-down algorithm-type’ method with isotonic regression for the estimation of a set of possible MTDs [6]. Furthermore, Wang and Ivanova have developed a three-parameter model-based method for which the parameters are estimated using Bayesian inference [7]. Mandrekar *et al.* have proposed incorporating both the toxicity and efficacy of each agent into the identification of an optimal dosing region for the combination using a continuation ratio model to separate each agent’s toxicity and efficacy curves [8, 9]. Fan *et al.* have proposed a ‘2 + 1 + 3’ algorithm-based dose-allocation scheme as well as the performance of two-dimensional isotonic regression to estimate the MTD [10]. Yin and Yuan have developed a Bayesian adaptive design based on latent 2×2 tables in which the combinations’ toxicity probabilities in the two-dimensional space are estimated using a Gumbel-type model [11]. Additionally, Yin and Yuan have extended their method by changing to a copula-type model to simulate the effect of two or more drugs in combination [12]. Bailey *et al.* have introduced a second agent as a covariate in a logistic model [13]. Recently, Wages *et al.* have considered a continual reassessment method (CRM) based approach considering different orderings with partial order between combinations. In this case, the MTD is estimated for the order associated with the highest model-selection criterion [14].

For this simulation study, we focused on designs dealing only with toxicity. We selected six designs, which were either algorithm-based or model-based, and whose statistical estimation methods and allocation rules differed. The aim of this work was to provide statisticians involved in dose-finding studies with tools to evaluate combinations in order to select the most suitable design according to clinical trial’s combination and indication. We selected the methods of Ivanova and Wang [6], Ivanova and Kim [18], Wang and Ivanova [7], Yin and Yuan [11, 12], and Wages *et al.* [14], which we consider representative of the methods that can be found in the literature. We had initially considered the selection of only one MTD at the end of the trial but then extended the comparison to the selection of multiple MTDs (Section 3).

2. Methods

2.1. Notations

Let there be a two-agent combination used in a phase I dose-finding clinical trial for which the dose-toxicity relationship is monotonic with respect to both dose levels. Let $D_{j,k}$ denote the dose level of a combination in which j refers to agent 1 ($j = 1, \dots, J$), and k to agent 2 ($k = 1, \dots, K$).

Toxicity refers to a dose-limiting toxicity (DLT), that is, an adverse event of grade 3 or higher. We represent the observation of a toxicity for patient i ($i = 1, \dots, N$) by a Bernoulli random variable y_i , equal to 1 if a DLT is observed for patient i and 0 otherwise. Let us assume that the combination dose $D_{j,k}$ is administered to $n_{j,k}$ patients and that we observe a total of $t_{j,k}$ DLTs at that combination dose

level. We then denote the observed proportion of DLTs by $\varphi_{j,k} = \frac{t_{j,k}}{n_{j,k}}$ and define $\pi_{j,k}$ as the toxicity probabilities.

The target (probability of) toxicity, θ , is fixed prior to trial onset as well as the initial guesses of toxicity (also called working model or skeleton). The combination X_i received by patient i ($i = 1, \dots, N$) can be seen as a random variable taking values $x_i \in \{D_{j,k}; j = 1, \dots, J; k = 1, \dots, K\}$. For simplicity purposes, the selected combination will be referred to as an MTD in order to maintain the same designation as in single-agent trials.

2.2. Up-and-down design for combinations [6]

Ivanova and Wang proposed an up-and-down design for two-agent combinations associated with an isotonic regression for the estimation of the MTD [6]. The aim of this method was to identify a set of MTDs for each dose of agent 2. Within this context, we proposed some modifications in order to ensure comparability with other methods. This two-dimensional approach was based on a non-parametric design, which is an extension of the Narayana group's design [15, 16]. If the last allocation was at combination $D_{j,k}$, the dose-allocation rule for the next combination is defined as follows: (i) $D_{j+1,k}$ if $\varphi_{j,k} < \theta$ and there was no toxicity observed in the last cohort; (ii) $D_{j-1,k}$ if $\varphi_{j,k} > \theta$ and there was at least one toxicity in the last cohort; (iii) $D_{j+1,k-1}$ if $\varphi_{j,k} < \theta$ and there was at least one toxicity in the last cohort; and (iv) $D_{j-1,k+1}$ if $\varphi_{j,k} > \theta$ and there was no toxicity in the last cohort.

The number of patients assigned to the lowest level of agent 2 is restricted to $\frac{N}{K}$ to enable the method to explore other levels of agent 2. Ivanova and Wang [6] proposed a 'start-up phase' in order to gather enough information before estimating $\varphi_{j,k}$'s. The start-up phase is conducted according to the following algorithm: (i) if, in the last cohort, no toxicity was observed, agent 1 is increased by one dose level; and (ii) if, in the last cohort, at least one toxicity was observed, agent 1 is decreased by two dose levels and agent 2 increased by one dose level. This process is repeated until all levels of agent 2 have been explored, and alternative combinations are proposed when reaching the boundary of the combination space.

When the overall sample size is reached, the estimate of the set of maximum tolerated combinations is calculated after using bivariate isotonic regression [17].

In order to ensure comparability between all of the methods presented in this manuscript, a decision rule with the selection of one MTD was proposed. The recommended combination at the end of the trial was the one with the toxicity probability closest to the target after isotonic regression; and, if there were several, then the one with the highest level of agent 2 was recommended. Moreover, the start-up phase was modified to avoid safety concerns. Indeed, we supposed that when combining two cytotoxic agents (and due to the potential synergistic effect in terms of toxicity between them), it is unreasonable to explore all levels of agent 2 during the start-up phase. Therefore, when toxicity was observed and the original rule was not possible, the start-up phase was stopped.

2.3. Up-and-down design using the T -statistic [18]

Using a newer approach proposed by Ivanova and Kim, a modification of the previous algorithm-based method can be implemented by replacing the Narayana design-based allocation rule with the T -statistic [18]. With the other parts of the method remaining the same, we defined the T -statistic at combination $D_{j,k}$ by the following:

$$T_{j,k} = \frac{\varphi_{j,k} - \theta}{\frac{s_{j,k}}{\sqrt{n_{j,k}}}}, \text{ where } s_{j,k}^2 = \frac{t_{j,k} - 2t_{j,k}\varphi_{j,k} + n_{j,k}\varphi_{j,k}^2}{n_{j,k} - 1}$$

Then, according to the recommendation on parameter values in [18], the dose-allocation rule would be as follows: (i) $D_{j+1,k}$ if $T_{j,k} \leq -1$, (ii) $D_{j-1,k}$ if $T_{j,k} \geq 1$, (iii) $D_{j+1,k-1}$ if $-1 < T_{j,k} \leq 0$, and (iv) $D_{j-1,k+1}$ if $0 < T_{j,k} < 1$.

2.4. Two-dimensional dose finding in discrete dose space [7]

Furthermore, Wang and Ivanova proposed a new two-dimensional model-based design of which the aim was to identify a set of MTDs for each dose of agent 2 [7]. As presented in the previous section, some

minor changes were proposed for this method in order to ensure comparability and respect clinical practice. The dose-combination model was defined as follows with an interaction term proposed by Wang and Ivanova:

$$\pi_{j,k} = 1 - (1 - p_j)^\alpha (1 - q_k)^\beta \exp(-\gamma \log(1 - p_j) \log(1 - q_k))$$

where $\alpha > 0$, $\beta > 0$, $\gamma > 0$, and p_j ($j = 1, \dots, J$), q_k ($k = 1, \dots, K$) are the working models for agents 1 and 2, respectively. The interaction term γ was introduced in the model to consider the possible synergistic effects.

After each cohort of patients, the estimation of $\hat{\pi}_{j,k}$ was updated using Monte Carlo method with exponential prior distributions centered in 1 for all parameters. At each step, the combination chosen to be allocated to the next cohort is the closest to the target belonging to $(D_{j+1,k}, D_{j,k+1}, D_{j-1,k+1}, D_{j-1,k}, D_{j,k-1}, D_{j+1,k-1}, D_{j,k})$.

For comparison purposes, the method was restricted to select only one MTD such that $D_{j,k}$ was the combination with a toxicity probability closest to the target: $(j, k) = \operatorname{argmin}_{j,k} |P(Y = 1 | D_{j,k}) - \theta|$ and among the $D_{j,k}$'s that have already been administered to patients, as proposed by Yin and Yuan [11, 12], without decreasing the performance of the method. Again, for the same reasons as previously outlined, the start-up phase was modified as detailed in Section 2.2.

2.5. Continual reassessment method for partial ordering [14]

Wages *et al.* proposed a dose-finding approach based on the CRM that considers orderings between combinations [14]. The ordering between agents is assumed to be monotonic and increases with the dose.

Assuming there are M possible ways to order combinations that are consistent with the non-decreasing assumption, let w_ℓ ($\ell = 1, \dots, J \times K$) be the working model (that is the initial guesses of toxicity in ascending order), and $\alpha_{i,\ell}$ the initial guess w corresponding to the position of the combination received by patient i ($i = 1, \dots, N$) in the ordering ℓ .

For each ordering $m = 1, \dots, M$, the dose-toxicity model is defined as a function of the dose and a parameter $a \in \mathbb{A} : \forall m, P(Y_i = 1 | X_i = x_i) = \pi_m(x_i, a)$, where the 'empiric' model $\pi_m(x_i, a) = \alpha_{i,m}^a$ was chosen by the authors. Following O'Quigley *et al.* [2], a prior probability distribution $g(a)$ for a was assigned. Let $\{p(1), \dots, p(M)\}$ denote the prior probability of each ordering representing their plausibility, where $\sum_m p(m) = 1$ and $\forall m, p(m) \geq 0$.

After the inclusion of I^{th} patient, for each ordering m , the posterior mean \hat{a}_m and the posterior probabilities of m are estimated by

$$\hat{a}_m = \frac{\int_{\mathbb{A}} a \cdot L_m(a|\text{data})g(a)da}{\int_{\mathbb{A}} L_m(a|\text{data})g(a)da} \quad \text{and} \quad \tilde{p}(m|\text{data}) = \frac{p(m) \int_{\mathbb{A}} L_m(a|\text{data})g(a)da}{\sum_{m=1}^M \left[p(m) \int_{\mathbb{A}} L_m(a|\text{data})g(a)da \right]}$$

where $L_m(a|\text{data})$ is a binomial likelihood.

The order h ($h = 1, \dots, M$) with the greatest posterior probability, $\tilde{p}(h|\text{data})$, is chosen for the next cohort. Nevertheless, when a uniform prior distribution is chosen for the ordering probabilities, as the trial proceeds, the difference between the posterior probabilities of m takes some time to differentiate. Therefore a start-up phase could be set up, where, for the first few patients, the ordering is sampled randomly with the weights of posterior probabilities. Once the order h has been chosen, the combination assigned to the next patient (or cohort) is the one closest to the target toxicity: $\operatorname{argmin}_\ell |\pi_h(w_\ell, \hat{a}_h) - \theta| = \operatorname{argmin}_\ell |w_\ell^{\hat{a}_h} - \theta|$ with $\ell = j \times k$. The original method recommended trial initiation at the combination assumed to be the MTD. Nevertheless, in our simulations, the first administered combination was the lowest for comparison purpose.

2.6. Dose-finding design based on copula regression [12]

Yin and Yuan proposed a Bayesian method using copula regression for combinations [12]. The authors have made the assumption that each single agent had already been evaluated in separate phase I trials.

As a consequence, physicians have a reasonable prior knowledge of the MTD of each drug alone. Let p_1, \dots, p_J and q_1, \dots, q_K be the prior toxicity probabilities of each dose level of agents 1 and 2 alone, respectively, and $\pi_{j,k}$ the probability of toxicity in combination ($D_{j,k}$).

A Clayton-copula regression type that enables expressing the joint toxicity probability of combination $D_{j,k}$ with marginal true probabilities of toxicity (p_j^α, q_k^β) was used:

$$\pi_{j,k} = 1 - \left((1 - p_j^\alpha)^{-\gamma} + (1 - q_k^\beta)^{-\gamma} - 1 \right)^{-\frac{1}{\gamma}}$$

where $\gamma, \alpha, \beta > 0$ are unknown parameters. The parameter γ characterizes the drug interaction effect, and α and β characterize the uncertainty on the initial guesses.

Let c_e and c_d , with $c_e + c_d > 1$, denote fixed probability cutoffs for dose escalation and de-escalation respectively that need to be calibrated through simulations studies. Prior distributions of model parameters are assumed to be independent with a prior distribution centered on 1 for α and β and a non-informative prior distribution for γ . Adaptive rejection Metropolis sampling within Gibbs sampling [19] was used to sample (α, β, γ) from the posterior joint distribution in order to calculate posterior estimates of $\pi_{j,k}$ and $P(\pi_{j,k} < \theta)$. In practice, the dose-allocation method is as follows: (i) if, at the current dose combination, $D_{j,k}$, $P(\pi_{j,k} < \theta) > c_e$, then the combination is escalated to an adjacent combination ($D_{j+1,k}, D_{j,k+1}, D_{j+1,k-1}, D_{j-1,k+1}$) with the probability of toxicity higher than the current value and closest to θ ; (ii) if, at the current dose combination, $D_{j,k}$, $P(\pi_{j,k} > \theta) > c_d$, then the combination is de-escalated to an adjacent combination ($D_{j-1,k}, D_{j,k-1}, D_{j+1,k-1}, D_{j-1,k+1}$) with the probability of toxicity lower than the current value and closest to θ ; and (iii) otherwise, the next cohort of patients continues to be treated at the current combination. Once the maximum sample size is reached, the combination associated with probability of toxicity that is closest to θ is selected as the MTD combination (from the dose tested on at least one cohort).

A start-up phase was proposed in order to gather enough information for estimating the $\pi_{j,k}$ where each agent's dose level is increased until at least one toxicity is observed while the other agent remains at its lowest level.

2.7. Dose-finding design based on latent contingency table [11]

The method proposed by Yin and Yuan in [11] is the same as that in [12] with another model for toxicity probability. A Gumbel model was chosen to model the probability of toxicity at combination $D_{j,k}$, given by

$$\pi_{j,k} = 1 - (1 - p_j^\alpha) (1 - q_k^\beta) \left[1 + p_j^\alpha q_k^\beta \frac{e^\gamma - 1}{e^\gamma + 1} \right]$$

3. Simulations

We simulated 2000 independent replications of phase I trials evaluating two-agent combination trials in which five dose levels for agent 1 and three dose levels for agent 2 were chosen. Eight scenarios were studied (Table I) with several number and locations of the MTDs in the combination space. The chosen scenarios seemed to cover a wide variety of underlying realities. The toxicity target was fixed at 0.3, and the overall sample size was 60. To ensure comparability, the cohort size was chosen equal to 3 for all methods, and no stopping rules were used. Because of practical concerns, each trial started at the lowest combination $D_{1,1}$.

For simplicity, the dose-finding methods are denoted in Section 3 as follows: (i) AISO for Ivanova and Wang [6], (ii) TSTAT for Ivanova and Kim [18], (iii) I2D for Wang and Ivanova [7], (iv) POCRM for Wages *et al.* [14], (v) BCOPULA for Yin and Yuan [12], and (vi) BGUMBEL for Yin and Yuan [11].

In order to be able to compare all dose-finding designs, modifications and assumptions were made (see Supporting information). All designs were optimized using the model average best-setting choices to improve the percentage of correct selections (PCS) when recommending one combination at the end of the trial. Indeed, we studied the influence of working models for each model-based design. Moreover, for each method with a start-up phase, we studied its influence. For example, for the POCRM, we studied the influence of the number of orderings retained for POCRM and the impact of those chosen on PCS. For I2D, we introduced the interaction term between the two agents suggested by Wang and Ivanova [7] and so on.

Table I. Toxicity scenarios for the two-agent combinations.

Agent 2	Agent 1									
	1	2	3	4	5	1	2	3	4	5
	Scenario 1					Scenario 2				
3	0.15	0.30	0.45	0.50	0.60	0.45	0.55	0.60	0.70	0.80
2	0.10	0.15	0.30	0.45	0.55	0.30	0.45	0.50	0.60	0.75
1	0.05	0.10	0.15	0.30	0.45	0.15	0.30	0.45	0.50	0.60
	Scenario 3					Scenario 4				
3	0.10	0.15	0.30	0.45	0.55	0.50	0.60	0.70	0.80	0.90
2	0.07	0.10	0.15	0.30	0.45	0.45	0.55	0.65	0.75	0.85
1	0.02	0.07	0.10	0.15	0.30	0.30	0.45	0.60	0.70	0.80
	Scenario 5					Scenario 6				
3	0.07	0.09	0.12	0.15	0.30	0.15	0.30	0.45	0.50	0.60
2	0.03	0.05	0.10	0.13	0.15	0.09	0.12	0.15	0.30	0.45
1	0.01	0.02	0.08	0.10	0.11	0.05	0.08	0.10	0.13	0.15
	Scenario 7					Scenario 8				
3	0.30	0.50	0.60	0.65	0.75	0.08	0.15	0.45	0.60	0.80
2	0.15	0.30	0.45	0.52	0.60	0.05	0.12	0.30	0.55	0.70
1	0.07	0.10	0.12	0.15	0.30	0.02	0.10	0.15	0.50	0.60
	Scenario 9					Scenario 10				
3	0.15	0.30	0.45	0.55	0.65	0.70	0.75	0.80	0.85	0.90
2	0.02	0.05	0.08	0.12	0.15	0.45	0.50	0.60	0.65	0.70
1	0.005	0.01	0.02	0.04	0.07	0.05	0.10	0.15	0.30	0.45

The MTD(s) combination are given in bold.

At the end of this optimization phase, in the simulation study, the marginal initial guesses of toxicities for agent 1, p_j , were chosen as (0.12, 0.2, 0.3, 0.4, 0.5), and for agent 2, q_k , as (0.2, 0.3, 0.4) for I2D, BCOPULA, and BGUMBEL using the ‘getprior’ function of the ‘dfcrm’ R package according to Lee and Cheung [20]. For BCOPULA and BGUMBEL, the dose-allocation thresholds were equal to $c_e = 0.8$, $c_d = 0.55$ and $c_e = 0.7$, $c_d = 0.55$, respectively, as proposed by Yin and Yuan [11, 12]. For POCRM, following Wages *et al.* [14], the number of possible orderings was restricted to 3 after a sensitivity analysis, and the working model was set up using the ‘getprior’ function with the length of indifference interval $\delta = 0.03$ and the initial guessed MTD $\ell = 13$ near the last combinations: (0.0001, 0.0006, 0.002, 0.005, 0.01, 0.02, 0.04, 0.06, 0.1, 0.14, 0.19, 0.24, 0.3, 0.36, 0.42). (Other working models were investigated; see Supporting information.)

At each simulated trial, we computed (i) the PCS at the end of the trial; (ii) the percentage of patients allocated at the true MTD(s) during the trial; (iii) the mean number of observed DLTs throughout the trial; and (iv) the mean number of patients allocated to each combination throughout the trial.

Designs were programmed using R version 2.13 [21] for AISO, TSTAT, I2D, and POCRM, and in C++ for BCOPULA and BGUMBEL.

3.1. Dealing with multiple MTDs

In this manuscript, we have proposed the recommendation of only one MTD at the end of the trial. In our case, we believe that the existence of one MTD for each row of agent 2 is not always true, but more than one MTD in the entire combination space is possible. Following this, we proposed some decision rules in order to identify at least one MTD at the end of the trial. We then evaluated its performance using the same scenarios as in the previous section. We first identified an MTD by level of agent 2; at the end of the trial, for $k = 1, \dots, K$, the MTD, $D_{j_k, k}$ is the combination closest to the target, as follows: $j_k = \operatorname{argmin}_j |P(Y = 1 | D_{j, k}) - \theta|$. Then we applied the following decision rules in order to identify MTDs that are too toxic or not toxic ‘enough’ by level of agent 2.

3.1.1. Decision rule for algorithm-based methods. The following decision rule was applied at the end of the trial: (i) if the combination selected to be the MTD $\in \{D_{1, k}, k = 1, \dots, K\}$ and $\pi_{1, k}^{\text{MTD}} - \theta > \tau_1$, then no combination was recommended on a row at the end of the trial; or (ii) if the combination selected

to be recommended $\in \{D_{J,k}, k = 1, \dots, K\}$ and $\theta - \pi_{J,k}^{MTD} > \tau_2$, then, once more, no combination was recommended.

3.1.2. Decision rule for model-based methods. The following decision rule was applied at the end of the trial: (i) if the combination selected to be recommended $\in \{D_{1,k}, k = 1, \dots, K\}$ and $P(\pi_{1,k}^{MTD} > \theta) > \tau_3$, then no combination was recommended on a row at the end of the trial; or (ii) if the combination selected to be recommended $\in \{D_{J,k}, k = 1, \dots, K\}$ and $P(\pi_{J,k}^{MTD} < \theta) > \tau_4$, then, once more, no combination was recommended. This rule was not applied to the POCRM as this method transforms a multidimensional combination space into an addition of several possible uni-dimensional orders.

In this simulation study, the thresholds were chosen as follows: $\tau_1 = \tau_2 = 0.15$, $\tau_3 = 0.90$, and $\tau_4 = 0.95$.

4. Results

4.1. Selection of one MTD at the end of the trial

Table II shows that the algorithm-based methods did not perform as well as the model-based ones. When comparing the performance of model-based methods, no design seemed to really stand out (Table II).

Scenarios 1 and 3 included three possible MTDs that were on one diagonal of the combination space; combinations $D_{2,3}$, $D_{3,2}$, and $D_{4,1}$ in scenario 1 and combinations $D_{3,3}$, $D_{4,2}$, and $D_{5,1}$ in scenario 3. For these scenarios, all model-based designs gave a high PCS (over 66%), whereas POCRM seemed to

Table II. Comparison of all dose-finding designs in terms of percentage of correct selection, percentage of patients allocated at the true MTD(s) during the trial, and mean number of observed DLTs throughout the trial when the aim is to select only one MTD.

	Scenario									
	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8	sc9	sc10
Percentage of correct selections										
AISO	46.9	57.9	57.4	36.4	67.7	39.0	50.7	27.0	24.8	31.8
TSTAT	55.0	52.5	62.3	36.8	67.1	45.3	46.5	30.9	20.8	27.0
I2D	68.0	73.7	66.9	89.7	83.7	37.2	41.9	50.4	5.1	13.0
POCRM	72.7	64.4	72.5	73.8	81.8	49.1	47.7	55.1	3.4	8.2
BCOPULA	66.2	71.8	71.7	84.1	78.1	30.7	49.6	43.5	5.0	16.3
BGUMBEL	67.1	72.5	68.4	87.5	77.9	33.6	48.0	49.5	6.0	8.6
CRM anti-diag1	73.7	74.8	71.9	84.9	80.0	71.4	73.2	84.3	0.0	0.0
CRM anti-diag2	73.7	74.8	71.9	84.9	80.9	75.1	75.4	83.9	0.0	0.0
Percentage of patient allocated at a true MTD(s) during the trial										
AISO	32.9	59.1	28.7	21.3	18.4	23.0	29.8	16.4	9.9	8.4
TSTAT	40.3	52.6	36.9	22.0	25.2	23.7	31.6	14.8	9.1	5.7
I2D	44.1	55.6	38.9	79.8	34.6	23.0	32.0	24.0	3.9	12.1
POCRM	46.8	39.6	51.6	57.4	66.1	28.8	34.1	28.5	3.1	8.8
BCOPULA	40.0	50.1	40.3	84.1	27.8	16.6	38.3	23.6	3.0	14.0
BGUMBEL	40.8	52.8	39.5	81.6	30.5	20.0	34.1	26.0	3.5	10.5
CRM anti-diag1	49.2	55.3	45.2	74.5	45.5	43.5	52.6	58.4	0.0	0.0
CRM anti-diag2	49.1	55.3	45.2	74.5	46.4	46.4	54.3	57.5	0.0	0.0
Mean number of observed DLTs all over the trial										
AISO	13.8	19.5	12.0	26.4	8.4	12.7	15.3	14.0	12.2	22.3
TSTAT	16.1	20.6	13.9	26.4	9.1	15.2	18.2	16.1	14.4	23.3
I2D	15.3	17.6	14.1	19.9	10.1	14.3	16.1	15.3	14.4	17.0
POCRM	20.1	22.8	18.2	23.3	14.4	17.7	19.9	20.5	17.4	22.5
BCOPULA	14.2	16.1	12.7	19.5	9.2	12.8	14.6	14.3	12.4	15.5
BGUMBEL	14.6	16.6	13.2	19.7	9.5	13.5	15.6	14.7	13.1	16.5
CRM anti-diag1	16.5	18.4	15.3	20.3	11.5	14.8	17.6	17.1	0.0	0.0
CRM anti-diag2	16.4	18.4	15.3	20.3	11.5	14.9	17.9	16.9	-	-

perform better in terms of PCS. In scenario 2, in which two possible MTDs were located in the lower part of the combination space, the highest PCS values were observed for I2D, BCOPULA, and BGUMBEL (over 70%). When the correct combination was in the lower ($D_{1,1}$) or higher ($D_{5,3}$) extremity of the combination space, as with scenarios 4 and 5, I2D performed better. In scenarios 6 and 7, when the true MTDs were randomly located in the combination space, the performance in terms of combination selection was low, less than 40% for most designs in scenario 6 and less than 50% in scenario 7. When there was only one true MTD and it was located in the middle of the combination space, the PCS was less than 55%, whichever the design. Finally, in scenarios 9 and 10, where the true MTD was unique and at the border of the combination space, the algorithm-based methods performed better than the model-based methods. For scenario 9, the PCS was above 20% for algorithm-based methods (AISO and TSTAT) but was always below 6% for model-based methods. This could be due to the way in which the combination space was explored: AISO and TSTAT provided better adjacent combination exploration owing to their dose-allocation method. Most PCS values remained, however, relatively low.

Table II shows that POCRM generated more DLTs than the other methods. It also tended to overtreat more patients than the other methods, and at higher combinations. In fact, the gain in PCS for POCRM

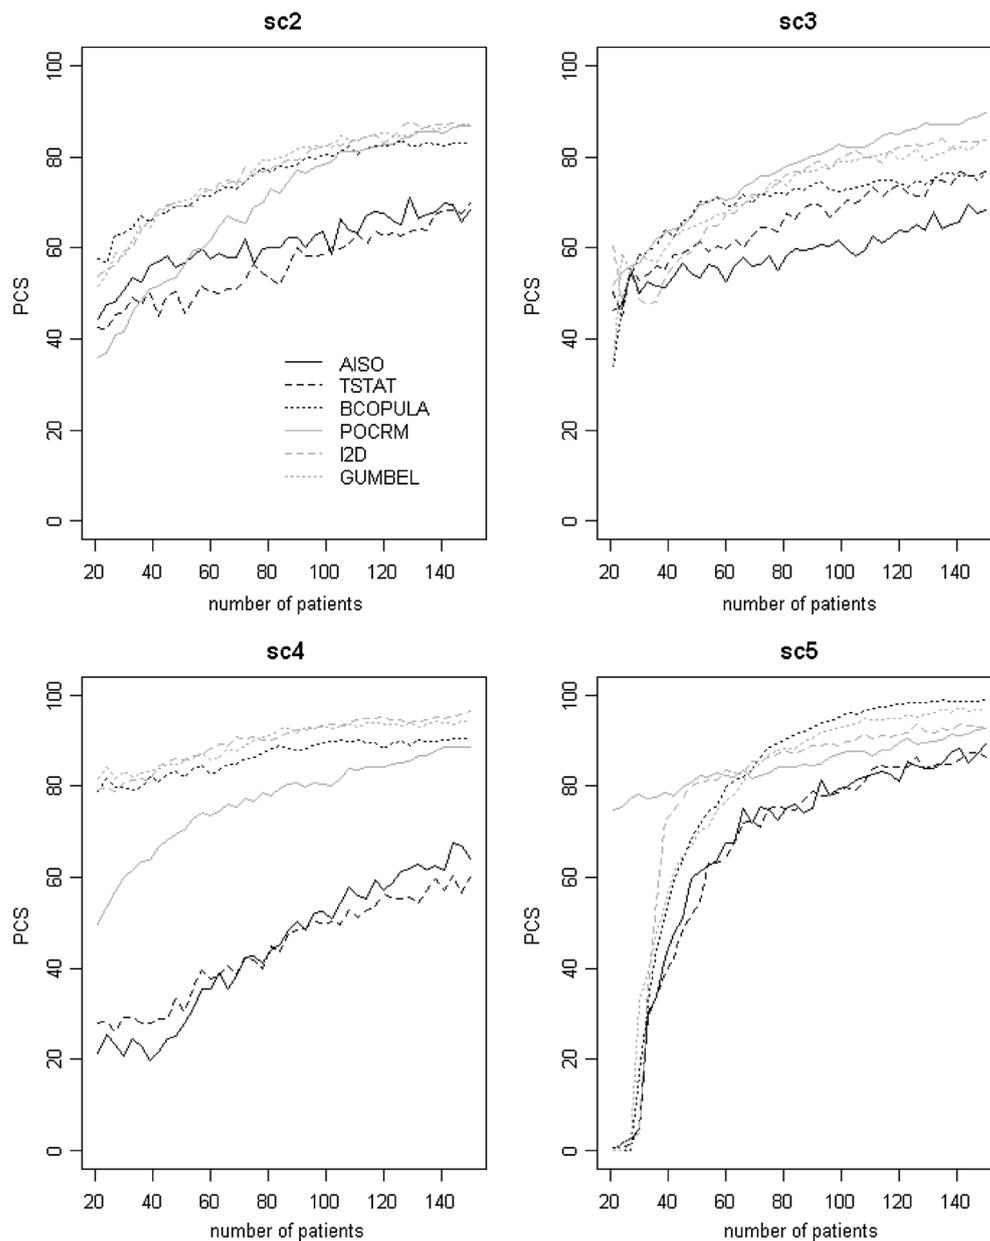


Figure 2. Convergence curves for scenarios 2–5.

when using a certain working model rather than others (see Supporting information) increased the mean number of DLTs. In the simulation study, we considered that this was acceptable, as the mean DLTs observed in all the scenarios was 19.7. This result is close to the expected number of DLTs (18), corresponding to a 0.3 toxicity combination in 60 patients. If investigators judged this possibility unacceptable, they would need to adjust the design parameterization prior to trial onset, or to consider a dose-allocation design with overdose control. The BCOPULA method gave similar PCS than BGUMBEL and had good properties in terms of mean DLTs number.

We then studied the convergence to the true MTDs while increasing the number of patients for all designs (Figure 2). We chose to show only four scenarios out of the eight presented in our simulation study. In scenario 2, all model-based designs were similar, but the algorithm-based design seemed to converge slowly. This finding was observed in all cases, where the convergence of AISO and TSTAT was slower than that of model-based methods. In scenario 3, the difference in PCS between the algorithm-based methods and the model-based methods tended to diminish with the increasing number of patients. In this scenario, POCRM, BGUMBEL, and I2D showed the best convergence, whereas in scenario 4, I2D and BGUMBEL approached nearly 95%. In scenario 5, BCOPULA and BGUMBEL had by far the best convergence and reached 90% very quickly. Nevertheless, in general, all methods (excepted BCOPULA and BGUMBEL in scenario 5) showed difficulties in attaining 100%, even with 300 patients. Overall, the convergence was rather slow.

4.2. Comparison with one-dimensional CRM

An important point is the contribution of multidimensional methods versus one-dimensional methods. As suggested during the review of this paper, we performed a one-dimensional CRM on a subset of combinations selected in an anti-diagonal of the dose-combination space where the toxicity probabilities order was known between combinations. We chose the following two different anti-diagonal paths chosen, as follows:

CRM anti-diag1: $D_{1,1} \rightarrow D_{1,2} \rightarrow D_{2,2} \rightarrow D_{3,2} \rightarrow D_{4,2} \rightarrow D_{5,2} \rightarrow D_{5,3}$

CRM anti-diag2: $D_{1,1} \rightarrow D_{2,1} \rightarrow D_{2,2} \rightarrow D_{3,2} \rightarrow D_{4,2} \rightarrow D_{4,3} \rightarrow D_{5,3}$

Using the ‘dfcrm’ package, we performed 2000 simulations on the scenarios corresponding to these anti-diagonals with restrictions to avoid skipping doses. The target toxicity, patient number, and cohort size were the same as for multidimensional methods. The working model was generated using the ‘get-prior’ function with an indifference interval $\delta = 0.05$, a initial guessed MTD at dose level 4 for a trial with seven doses. For scenarios 1–5, where at least one of the true MTD was included in anti-diagonals, PCSs were similar between multidimensional model-based methods and CRM. For scenarios 6–8, where the true MTDs were not located on the same diagonal, CRM on a reduced ordered subset of combinations containing at least one MTD had clearly higher performances than multidimensional designs. But in practice, the true MTD(s) is (are) not necessarily contained in the chosen anti-diagonal of the ordered

Table III. Comparison of AISO, TSTAT, I2D, BCOPULA, and BGUMBEL designs in terms of percentage of correct combination selection for each level of agent 2 when selecting multiple MTDs.

	Scenario 1			Scenario 2		Scenario 3			Scenario 4
	D _{4,1}	D _{3,2}	D _{2,3}	D _{2,1}	D _{1,2}	D _{5,1}	D _{4,2}	D _{3,3}	D _{1,1}
AISO	49.0	55.5	49.3	49.6	71.2	44.4	54.1	49.5	51.8
TSTAT	46.1	54.9	61.5	46.4	64.8	44.6	56.8	64.4	57.5
I2D	55.6	71.3	63.5	76.9	84.6	75.4	64.6	64.5	90.4
BCOPULA	41.0	47.9	35.2	70.9	82.5	47.0	55.4	33.8	73.2
BGUMBEL	46.1	61.3	38.5	76.5	83.2	40.9	67.0	41.2	74.5
	Scenario 5	Scenario 6		Scenario 7			Scenario 8	Scenario 9	Scenario 10
	D _{5,3}	D _{4,2}	D _{2,3}	D _{5,1}	D _{2,2}	D _{1,3}	D _{3,2}	D _{2,3}	D _{4,1}
AISO	67.3	47.6	42.2	34.4	47.8	41.8	76.3	37.7	46.9
TSTAT	69.5	49.9	48.4	33.2	40.8	62.5	75.4	40.3	39.8
I2D	77.9	60.9	32.4	18.1	42.6	53.9	89.3	11.8	23.4
BCOPULA	80.7	23.8	20.5	23.4	33.1	63.4	86.2	5.9	31.4
BGUMBEL	83.2	35.6	25.0	7.8	41.4	64.0	91.5	8.2	19.6

combinations retained, which was the case of both scenarios 9 and 10. For these scenarios, the CRM (anti-diad1 and anti-diag2) could obviously never select a true MTD as it was not contained in the chosen path. In this case, algorithm-based multidimensional methods performed better than model-based ones, even if PCS remained quite low.

4.3. Selection of multiple MTDs at the end of the trial

When selecting one MTD per level of agent 2 (Table III), PCSs of all methods were good on each row for scenarios 1–5 and 8 (higher than 40% in all cases and up to 91.5%). For scenario 6, the algorithm-based methods (AISO and TSTAT) and I2D performed well, whereas BCOPULA and BGUMBEL had rather low PCSs (between 20% and 35%). For scenario 7, $D_{2,2}$ and $D_{1,3}$ were well identified by all designs, but the PCSs for $D_{5,1}$ were lower for model-based methods.

5. Discussion

The aim of this manuscript was to compare several dose-finding designs for cytotoxic combination studies. Based on this simulation study, model-based methods seemed to perform better than algorithm-based methods in terms of the percentage of correct combination selections (PCSs) when targeting a single MTD at the end of the trial. In general, the model-based methods gave a high PCSs in this case, and there was no major difference between the model-based methods compared. When one MTD per row was targeted, algorithm-based methods performed better than model-based methods but with low PCS.

For comparison purposes, several choices were made, which merit discussion. According to the combination dimensional space, we arbitrarily fixed the sample size at 60, as in Yin and Yuan [11, 12]. In this study, we chose five dose levels of agent 1 and three dose levels of agent 2, which resulted in 15 possible combinations to evaluate. Nevertheless, when using a different dimensional space ($J \times K$), further investigations need to be carried out to find the optimal sample size for each method. In practice, it seems unreasonable to have such a large number of available combinations to evaluate, and only a subset of the dimensional space could be relevant. For this reason, we decided to compare the methods on a more realistic basis. Therefore, we chose 10 scenarios on a 5×2 dimensional space and performed 2000 simulations of trials with 40 patients (data not shown). As in our manuscript, all model-based designs performed well.

Some authors have made the assumption that using one MTD for each level of agent 2 is possible when exploring a large number of combinations [6, 7]. We thus proposed decision rules designed to detect when at least one MTD existed in the combination space. These decision rules were implemented at the end of the trial and were found to maintain the performance of the designs. In this case, how should the most appropriate combination for further investigation be chosen? Phase II trials can study several combinations, and if they require the selection of a unique combination, other criteria such as efficacy or pharmacokinetics should be taken into account in the decision process. Indeed, when two cytotoxic agents are combined, the resulting pharmacokinetic (PK) profiles of two MTDs are not necessarily similar. In this case, the investigators could base their final decision on the maximization of exposure, or on the maximization of an efficacy surrogate.

Some issues are raised by the design modifications proposed in this paper. Some methods are designed to select only one MTD. For instance, BCOPULA and BGUMBEL have a conservative allocation algorithm that explores a restricted subset of the combination space and focuses on one combination when it is estimated to be the correct one. In these methods, patients are often allocated to one or few combinations, and the other combinations are allocated to very few or to no patients. As a result, the estimation in a row of agent 2 can be poor. Moreover, for decision rules, we decided to keep the same tau values ($\tau_1 = \tau_2 = 0.15$, $\tau_3 = 0.90$, and $\tau_4 = 0.95$) for all of the designs. But some designs could have performed better if we had calibrated these values specifically.

The partial ordering method (POCRM) [14] is based on determining the most appropriate combination ordering in terms of toxicity, from a set of possible orderings. Nevertheless, the number of possible orderings increases with the combination space. In our simulation study, we restricted the choice to three reasonable orderings, as in Wages *et al.* [14] (see Supporting information). It should be noted that the method does not contain a ‘non-skipping’ rule and that in theory the combination allocated to the next cohort can ‘skip’ more than one combination (that is, selecting a combination, which is not in the immediate adjacent space of the current combination). Especially, if investigators necessarily

wish to begin the trial at the lowest combination, as we did in our simulation study, and no toxicity is observed, then the next combination will gravitate towards the initial guessed MTD. This will cause a huge skipping depending on the ordering selected and on the working model. It is clear that in practice this should not be allowed. As most real clinical trials begin at the lowest dose, we think that this method should be modified to allow initiating the trial at the lowest combination while adding ‘non-skipping’ rules or a start-up phase. For instance, as in ‘classical’ one-dimensional CRM, after selecting the ordering with the highest posterior probability, the method could restrict the dose allocation to combinations, which are next to the current one in the most probable ordering. Another possibility could be, if the current combination is $D_{j,k}$, to restrict the next combination when escalating to combinations $D_{j+1,k}$ and $D_{j,k+1}$ or to implement one of the starting phase proposed by Yin and Yuan [12] or by Wang and Ivanova [7].

During the review of this manuscript, Wages and Conaway [22] have published a paper proposing some guidelines for the POCRM. In their paper, they have suggested to place the initial guessed MTD at the middle of the working model to ensure that there is enough spacing both below and above this dose. In our simulation study, according to our sensitivity analyses, we have placed the initial guessed MTD near the third quartile of the dose range. As published by O’Quigley and Zohar [23], there is no sharp answer about what is the definition of a reasonable against a non-reasonable working model, although it may well be the notion of robustness itself. In the Supporting information, we have tried to point out how a non-reasonable or mis-specified choice can dramatically lessen the performance of the method (this choice was not robust for all scenarios; see Supporting information) [24]. This is why in this manuscript our choice was driven by this finding; thus we have selected a reasonable working model as it has shown to give good performances on average for all scenarios. Another important modification that we have added to the POCRM is the recommendation to start at the lowest dose. We have based our decision on common practice in phase I for a single agent or a combination of agents. This modification has shown to have equal performance than if the POCRM started at the initial guessed MTD (data not shown).

Another important issue relates to the performance of multidimensional methods versus one-dimensional methods. As suggested during the review of this paper, we performed a basic CRM on a subset of combinations selected in an anti-diagonal of the dose-combination space. When the MTD was included in the anti-diagonal, the one-dimensional basic CRM worked as well or better than any multidimensional method. This finding points out that if the MTD exists in the selected anti-diagonal, a one-dimensional method is preferable to a more complex one. In practice, the entire combination space is not often studied in combination trials; it can increase the number of combinations to be evaluated, and ‘3 + 3’ dose-allocation rules, which are still used by investigators, are not valid for such trials. Using one-dimensional approaches involves a choice by the investigators in the determination of the combinations to study, and this can be a difficult question. The most important issue will then be whether the chosen subset of combinations contains an MTD. If it does, a one-dimensional method would perform better than multidimensional ones.

In our comparative simulation study, none of the model-based designs gave consistently better results than the others. Each method requires several choices prior to trial beginning, such as the choice of the working model, of the start-up phase, and of the prior distributions. According to our simulation results (see also Supporting information), it seemed that some choices can tumble the performance of a design. The issue of using a single MTD or multiple MTDs when evaluating a large combination space is challenging. Statisticians should propose combination methods that could identify the presence of one or more MTDs in the combination space, in their assumption, and in the dose-allocation process. These methods should also identify at which levels of agent 2 MTDs are located. Statisticians and investigators should be aware of the pros and cons of these designs in planning future trials. Our work was to enlighten multidimensional methods by comparing them using the same scenarios and the same (or very close) features.

Acknowledgements

We would like to thank Drs Ying Yuan and Raphael Porcher as well as reviewer 1 and the associate editor for their advice and constructive comments that have served to substantially improve the manuscript. This work was partially funded by grants from the ANRT (Association Nationale de la Recherche et de la Technologie), ‘Laboratoires Servier’ (CIFRE number 2011/0900), and Denis Diderot University-Paris 7 (‘Aide à la mobilité internationale 2012’).

References

1. Storer B. Design and analysis of phase I clinical trials. *Biometrics* 1989; **45**:925–37.
2. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics* 1990; **46**:33–48.
3. Faries D. Practical modifications of the continual reassessment method for phase I cancer clinical trials. *Journal of Biopharmaceutical Statistics* 1994; **4**:147–64.
4. Korn E, Midthune D, Chen T, Rubinstein L, Christian M, Simon R. A comparison of two phase I trial designs. *Statistics in Medicine* 1994; **13**:1799–806.
5. Chevret S. *Statistical Methods for Dose-Finding Experiments*, Statistics in Practice. John Wiley and Sons Ltd.: Chichester, 2006.
6. Ivanova A, Wang K. A non-parametric approach to the design and analysis of two-dimensional dose-finding trials. *Statistics in Medicine* 2004; **23**:1861–1870.
7. Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics* 2005; **61**:217–222.
8. Mandrekar S, Cui Y, Sargent D. An adaptive phase I design for identifying a biologically optimal dose for dual agent drug combinations. *Statistics in Medicine* May 2007; **26**:2317–2330.
9. Mandrekar S, Qin R, Sargent D. Model-based phase I designs incorporating toxicity and efficacy for single and dual agent drug combinations: methods and challenges. *Statistics in Medicine* May 2010; **29**:1077–1083.
10. Fan S, Venook A, Lu Y. Design issues in dose-finding phase I trials for combinations of two agents. *Journal of Biopharmaceutical Statistics* 2009; **19**:509–523.
11. Yin G., Yuan Y. A latent contingency table approach to dose finding for combinations of two agents. *Biometrics* 2009; **65**:866–875.
12. Yin G, Yuan Y. Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society* 2009; **58**:211–224.
13. Bailey S, Neuenschwander B, Laird G, Branson M. A Bayesian case study in oncology phase I combination dose-finding using logistic regression with covariates. *Journal of Biopharmaceutical Statistics* 2009; **19**:469–484.
14. Wages N, Conaway M, O'Quigley J. Continual reassessment method for partial ordering. *Biometrics* 2011; **67**:1555–63.
15. Narayana T. Sequential procedures in the probit analysis. *Dissertation*, The University of North Carolina, Chapel Hill, 1953.
16. Ivanova A, Montazer-Haghigh A, Mohanty S, Durham S. Improved up-and-down designs for phase I trials. *Statistics in Medicine* 2003; **22**:69–82.
17. Dykstra R, Robertson T. An algorithm for isotonic regression for two or more independent variables. *Annals of Statistics* 1982; **10**:708–716.
18. Ivanova A, Kim S. Dose finding for continuous and ordinal outcomes with a monotone objective function: a unified approach. *Biometrics* Mar 2009; **65**(1):307–315.
19. Gilks W, Best N, Tan K. Adaptive rejection metropolis sampling within Gibbs sampling. *Journal of the Royal Statistical Society* 1995; **44**:455–472.
20. Lee S, Cheung Y. Model calibration in the continual reassessment method. *Clinical Trials* 2009; **6**:227–38.
21. R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing: Vienna, Austria, 2011. <http://www.R-project.org/>, ISBN 3-900051-07-0.
22. Wages N, Conaway M. Specifications of a continual reassessment method design for phase I trials of combined drugs. *Pharmaceutical Statistics* 2013; **12**(4):217–224.
23. O'Quigley J, Zohar S. Retrospective robustness of the continual reassessment method. *Journal of Biopharmaceutical Statistics* Sep 2010; **20**(5):1013–1025.
24. Neuenschwander B, Branson M, Gsponer T. Critical aspects of the Bayesian approach to phase I cancer trials. *Statistics in Medicine* Jun 2008; **27**(13):2420–2439.

Supporting information

Additional supporting information may be found in the online version of this article at the publisher's web site.

1 Additional material - Optimisation

Table 1: Working models retained to study their influence on PCS for I2D, BCOPULA and BGUMBEL. These working models were constructed according to Lee and Cheung [1].

	working model for agent 1					for agent 2		
WM1	0.025	0.06	0.12	0.20	0.30	0.12	0.20	0.30
WM2	0.12	0.20	0.30	0.40	0.50	0.20	0.30	0.40
WM3	0.30	0.40	0.50	0.59	0.67	0.30	0.40	0.50

Table 2: PCSs for I2D, BCOPULA and BGUMBEL according to the three working models.

		Percentage of correct selections (PCS)								
		scenario	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8
I2D	WM1	64.9	67.6	59.7	73.5	88.6	43.1	40.0	48.5	
	WM2	61.6	66.2	58.0	93.6	90.8	35.4	43.3	42.0	
	WM3	57.3	48.6	51.4	97.2	92.7	33.6	41.0	34.4	
BCOPULA	WM1	59.0	71.1	64.8	81.8	89.8	30.2	50.9	39.3	
	WM2	66.2	72.4	71.7	84.5	78.1	30.7	49.4	43.5	
	WM3	72.1	72.1	65.8	89.5	11.8	40.4	36.7	44.2	
BGUMBEL	WM1	60.9	69.8	63.7	84.0	89.0	33.2	45.6	44.8	
	WM2	67.1	72.5	68.4	87.5	77.9	33.6	48.0	49.5	
	WM3	73.9	71.2	67.4	91.8	11.4	41.0	49.8	54.6	

1.1 I2D

1.1.1 Influence of the working model

Table 1 gives three working models, WM1, WM2 and WM3, in which the target toxicity was shifted from the end to the beginning of the dose levels. All three working models were constructed according to Lee and Cheung [1]. The PCSs of the eight scenarios for each working model are given in Table 2.

The PCSs for WM3 were always lower than for the other working models, except for scenarios 4 and 5 where the true MTD lay in the extremities (Table 2). It is of note that the PCSs for scenarios 4 and 5 were already very high, irrespective of the working model.

It was not easy to make a choice between WM1 and WM2. We chose WM2 for the manuscript; this choice can be discussed when considering WM1 or WM2.

1.1.2 Interaction term

The possible synergistic effect was taken into account in the model by introducing an interaction term $\gamma > 0$ as proposed by Wang and Ivanova [2]. The model was then defined as follows:

$$\pi_{j,k} = 1 - (1 - p_j)^\alpha (1 - q_k)^\beta \exp(-\gamma \log(1 - p_j) \log(1 - q_k)).$$

The joint distribution, g , of parameters (α, β, γ) is then the product of three independent exponential distributions with mean equal to 1.

A simulation study was performed in order to detect whether the interaction term could improve our results (Table 3). In most scenarios, the interaction term improved the PCSs. As a result, the interaction term was retained.

Table 3: PSCs for the I2D design according to the presence or absence of the interaction term γ in the model.

scenario	Percentage of correct selections (PCS)							
	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8
no interaction	61.6	66.2	58.0	93.6	90.8	35.4	43.3	42.0
interaction	67.4	68.2	66.6	94.2	84.7	32.6	45.0	44.6

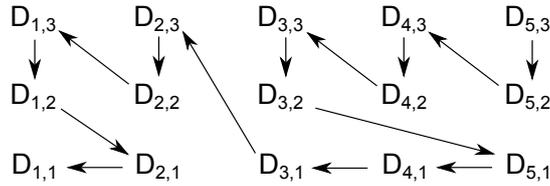
1.2 POCRM

1.2.1 Influence of the chosen orderings

The number of possible orderings increases with the number of combinations. In practice, it is more reasonable to choose a restricted number of orderings in the clinical trial. First we studied the influence of a restricted number of orderings with specific shapes (i.e., repeated pattern). Then we randomly selected three orderings from all the possible orderings consistent with the partial order (results not shown).

In general, and except for variations, choosing orderings with specific shapes provides high PCSs. When the orderings are chosen randomly (like Figure 1), PCSs can fall very low in some scenarios, for instance, when there is only one MTD in the combination space and in the orderings retained, the toxicity probabilities before the true MTD are high. Therefore, if these too toxic combinations are tested, the method will de-escalate and miss the true MTD.

Figure 1: Ordering causing a fall in PCS in scenario 8 with random orderings for POCRM.



1.2.2 Influence of the number of orderings

After studying the influence of the retained ordering, we chose to study the influence of their number. In this context, the PCSs with 3, 7, 22 or 200 orderings were considered. The first three orderings were those chosen for the manuscript. It should be noted that the number of possible orderings was equal to 6006.

Table 4: PCSs for POCRM using working model 2 with an increase in the number of orderings from 3 to 200.

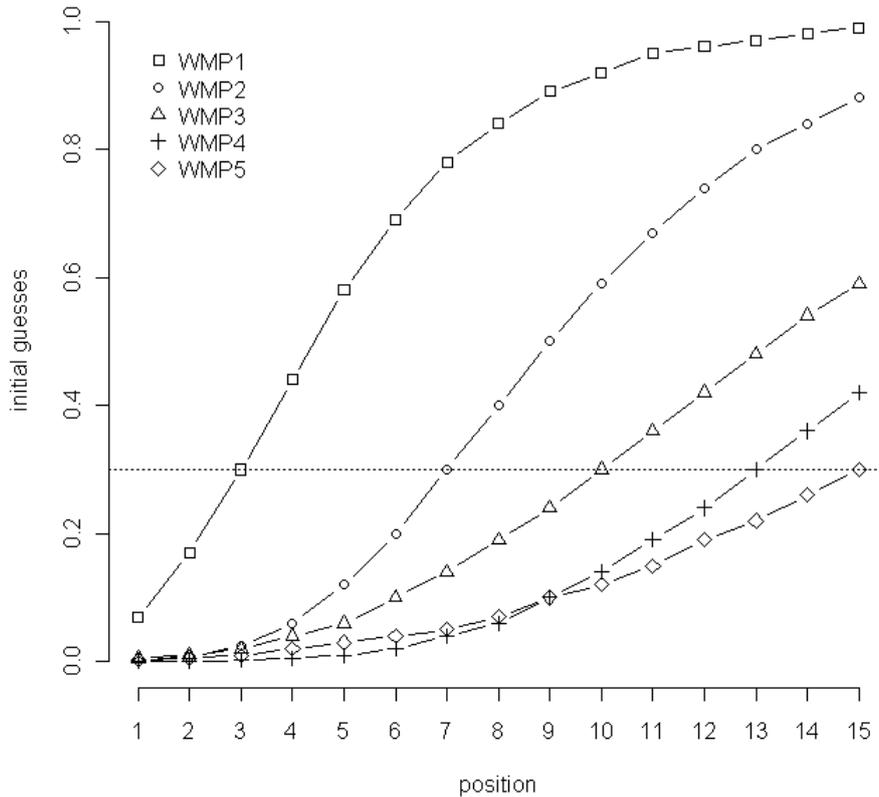
scenario	Percentage of correct selections (PCS)							
	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8
number of orderings including those with a specific shape								
3	78.3	60.0	73.9	65.6	15.8	57.4	49.1	58.6
7	69.8	56.7	72.3	62.2	13.7	61.3	65.8	55.1
22	69.1	57.4	73.0	64.5	10.9	55.4	72.5	51.9
200	70.1	56.8	72.9	62.5	7.5	57.1	73.0	49.7

In general, increasing the number of orderings did not seem to improve the performance of the design, or even reduce it, except for scenarios 6 and 7. Indeed, these scenarios were more peculiar (the MTDs were not on the same diagonal of the combination space, see Table 4), and the retained orderings did not correspond to it. Therefore the higher the number of orderings, the easier it was to find an ordering that can fit this scenario.

1.2.3 Influence of the working model (WMP)

As for the standard CRM, the influence of the working model should be taken into consideration. It was noticed that the PCS for scenario 5 was always very low, and our presumption was that it could be due to the choice of the working model. According to Lee and Cheung’s skeleton construction [1], several working models were defined from which we chose to put the initial guessed MTD at different positions: 3, 7, 10, 13 and 15 (Figure 2). Is it important to note that due to the high number of combinations studied, setting the initial guessed MTD at an extremity could result in a misspecified working model as for WMP1 [3], nevertheless the aim of this analysis was also to study the implications of such choice. In these simulations, we chose the same three orderings as in the manuscript.

Figure 2: Five different working models used in the simulation study for POCRM.



From our simulation study, we noticed that the lower the position of the initial guessed MTD, the lower the PCS for this working model, when the true MTDs were in the middle or at the end of the combination space (cf. values in bold and underlined for smaller decreases in Table 5). Indeed, the position of the initial guessed MTD still remained highly influential. It seemed that the higher the number of combinations, the more the CRM has some difficulty in identifying the MTDs according to the initial guessed position of the MTD in the working model. Increasing the number of patients in

some cases did not improve the results. For instance, when the number of patients was increased from 60 to 300, the PCS of scenario 5 for POCRM was equal to 72.2%, whereas it was equal to 82.7% with a working model for which the initial guessed MTD position was at the 13th position with 60 patients. It is to be noted that, on the contrary, if the initial guessed MTD position is at the end, the method has no difficulty in decreasing.

Table 5: Percentage of correct selection of POCRM using five different working models.

scenario	Percentage of correct selections (PCS)							
	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8
WMP1	<u>44.5</u>	81.8	3.9	71.3	0.0	36.7	76.9	19.0
WMP2	78.3	60.0	73.9	65.6	15.8	57.4	49.1	58.6
WMP3	74.5	63.1	72.1	75.8	78.2	48.6	49.0	51.5
WMP4	72.2	60.5	72.5	72.6	82.7	49.1	47.7	50.5
WMP5	73.5	65.2	70.8	77.8	83.6	45.0	37.6	49.6

WMP4 and WMP5 were both efficient. According to the PCS, WMP5 could be considered as better, but the number of patients exposed to overdose was higher than for WMP4. Indeed, at the beginning of the trial, the initial guessed MTD position had a high influence on the dose allocation process. Therefore, if the true MTD is located in the lower combinations, patients can easily be overdosed. A good compromise between PCS and ethical criteria was to choose WMP4 for our manuscript.

1.3 BCOPULA

1.3.1 Influence of the working model

We studied the influence of the working model on the performance of the design. According to Table 2, when the initial guessed MTD is chosen in the first combinations, the method would have difficulty selecting the true MTD in scenario 5 (PCS equal to 11.8%). That is why WM3 was not retained for its very low PCS in this scenario. In general, the results decreased in all scenarios, except in scenario 5, when the initial guessed MTD position was located at the end, whereas in scenario 5 it increased substantially. Therefore, both WM1 and WM2 could be chosen. We chose WM2 for our manuscript due to PCSs that seemed quite balanced between all the scenarios.

2 Additional material - Combination allocation

In scenarios 1 and 3, the distribution of allocated dose levels around the MTDs (see Tables 6 and 7) seemed to be spread unevenly. For example, for POCRM, more patients were allocated to the higher dose levels, and the opposite was the case for I2D; that is, more patients were allocated to the lower dose levels. In scenario 2, the methods with the highest PCS values have more than 50% patients allocated to the MTDs. For scenario 5, the observed difference, in terms of allocated patients at the true MTD during the trial (Table 6) between POCRM (66.1%) and BCOPULA (27.8%) or BGUMBEL (30.5%), was due to the dose-allocation rule. Indeed, in POCRM, the best ordering was chosen at each step with no restriction on combinations skipping, whereas BCOPULA or BGUMBEL included a skipping restriction, as the next allocated combination depended on the location of the current one.

3 Additional Material - Multiple MTDs

In Tables 8 and 9 are given the PCS when selecting multiple MTDs and using decision rules described in simulation section.

References

- [1] Lee S, Cheung YK. Model calibration in the continual reassessment method. *Clin Trials* 2009; **6**:227–38.
- [2] Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics* 2005; **61**:217–222.
- [3] O’Quigley J, Zohar S. Retrospective robustness of the continual reassessment method. *J Biopharm Stat Sep* 2010; **20**(5):1013–1025.

Table 6: Mean number of patients allocated to each combination all along the trial ($N = 60$) for AISO, TSTAT and I2D designs.

	AISO			TSTAT			I2D								
sc 1	4.3	5.4	1.9	0.3	0.1	3.0	11.9	6.6	1.3	0.2	2.5	7.7	2.6	0.3	0.0
	3.9	8.7	9.0	3.9	0.7	1.3	4.2	7.7	3.5	0.7	1.3	4.7	9.3	2.6	0.5
	3.9	5.0	5.7	5.3	1.9	4.1	4.5	5.0	4.5	1.6	3.7	4.2	7.1	9.5	3.3
sc 2	1.0	0.2	0.0	0.0	0.0	4.4	2.0	0.4	0.0	0.0	3.4	0.4	0.0	0.0	0.0
	27.5	7.7	1.1	0.2	0.1	21.8	7.3	1.3	0.5	0.1	15.2	4.5	0.2	0.0	0.0
	8.2	7.9	3.3	1.6	1.1	7.9	9.8	2.5	1.0	1.0	14.8	18.2	2.4	0.4	0.0
sc 3	1.8	4.8	5.3	2.1	0.6	0.9	4.1	11.5	5.3	1.1	0.9	3.0	7.1	3.3	0.9
	1.3	3.4	7.7	8.5	3.1	0.8	1.8	4.8	7.3	2.5	0.7	1.7	5.2	8.3	4.2
	3.3	4.1	4.6	5.9	3.4	3.4	4.2	4.2	4.6	3.3	3.2	3.6	3.3	5.9	8.0
sc 4	0.1	0.0	0.0	0.0	0.0	1.2	0.4	0.0	0.0	0.0	0.2	0.0	0.0	0.0	0.0
	35.5	2.0	0.1	0.0	0.0	32.9	2.4	0.5	0.1	0.0	4.2	0.4	0.0	0.0	0.0
	12.8	5.0	1.9	1.3	1.2	13.2	5.2	1.5	1.2	1.2	47.9	6.7	0.2	0.0	0.0
sc 5	0.7	1.6	3.9	6.8	11.0	0.7	1.4	3.3	6.0	15.1	0.7	1.3	2.6	4.5	20.8
	0.6	1.5	3.6	5.1	5.9	0.6	1.4	3.2	4.1	5.1	0.6	1.3	2.8	2.6	5.7
	3.1	3.4	4.1	4.4	4.3	3.1	3.3	4.2	4.1	4.2	3.0	3.1	3.0	2.7	4.5
sc 6	2.9	5.3	3.7	1.0	0.5	1.9	8.7	9.2	3.5	0.8	1.5	5.1	6.2	2.2	0.7
	2.1	3.8	7.7	8.5	3.7	0.9	2.2	5.0	5.5	2.7	1.0	2.5	5.5	8.7	4.1
	3.8	4.4	4.4	4.7	3.4	3.9	4.3	4.1	4.1	3.1	3.6	3.8	3.4	4.4	6.6
sc 7	4.0	2.1	0.4	0.1	0.0	8.0	8.5	2.6	0.5	0.0	5.6	4.2	0.7	0.1	0.0
	10.4	11.3	6.4	2.9	0.9	3.8	8.6	4.6	2.2	0.6	3.5	8.6	6.6	2.2	0.5
	4.2	4.9	4.7	4.9	2.6	4.4	4.9	4.8	3.6	2.3	4.1	4.9	6.2	7.2	5.0
sc 8	3.1	6.4	3.8	0.5	0.0	1.7	10.2	9.9	0.9	0.0	1.6	7.0	2.3	0.2	0.0
	2.4	9.0	9.8	3.0	0.2	1.1	4.4	8.9	3.1	0.2	1.1	4.0	14.4	1.4	0.0
	3.5	5.2	6.9	4.9	1.2	3.4	4.6	5.8	4.4	1.2	3.3	4.1	10.0	8.8	1.1

In bold are given the MTD(s).

Table 7: Mean number of patients allocated to each combination all along the trial ($N = 60$) for POCRM, BCOPULA and BGUMBEL designs.

	POCRM			BCOPULA			BGUMBEL								
sc 1	1.7	9.0	6.1	3.5	0.7	8.1	8.4	2.9	0.1	0.0	7.4	8.1	3.6	0.3	0.0
	0.4	1.3	8.9	4.6	2.6	4.3	2.7	7.0	2.6	0.1	4.0	3.3	9.5	3.0	0.2
	3.0	0.2	2.2	10.2	5.0	3.8	3.5	5.7	8.6	2.1	3.6	3.5	5.2	6.9	1.4
sc 2	6.4	2.5	1.0	1.5	0.0	3.0	0.1	0.0	0.0	0.0	1.9	0.1	0.0	0.0	0.0
	11.6	5.4	1.6	0.4	0.7	16.5	3.2	0.1	0.0	0.0	13.2	4.8	0.2	0.0	0.0
	6.9	12.2	6.2	2.4	0.7	21.1	13.6	2.1	0.2	0.0	18.9	18.5	2.0	0.2	0.0
sc 3	0.2	2.3	9.4	7.0	1.6	4.9	3.2	6.3	1.7	0.3	4.2	3.3	7.3	3.4	0.4
	0.1	0.2	3.8	12.6	7.3	3.3	0.4	2.6	8.1	2.7	3.2	0.6	3.0	9.8	2.7
	3.0	0.0	0.2	2.6	9.0	3.2	3.2	3.2	7.1	9.8	3.2	3.2	2.9	6.3	6.6
sc 4	1.6	0.9	0.3	0.7	0.0	0.7	0.0	0.0	0.0	0.0	0.6	0.0	0.0	0.0	0.0
	7.8	1.5	0.3	0.0	0.3	4.9	0.1	0.0	0.0	0.0	4.4	0.2	0.0	0.0	0.0
	34.5	8.8	2.2	0.5	0.1	50.5	3.5	0.2	0.0	0.0	49.0	5.4	0.2	0.0	0.0
sc 5	0.0	0.1	1.2	7.9	39.7	4.0	1.1	2.0	3.3	16.7	3.7	1.1	1.8	3.7	18.3
	0.0	0.1	0.3	0.5	5.7	3.1	0.1	0.1	0.9	7.8	3.1	0.1	0.3	1.8	8.7
	3.0	0.0	0.0	0.2	0.6	3.0	3.0	3.1	3.9	7.7	3.0	3.0	2.9	3.3	5.3
sc 6	1.0	5.2	5.7	5.1	1.0	7.4	6.7	2.9	1.0	0.7	6.1	6.5	4.5	1.9	0.6
	0.3	0.8	7.3	12.1	8.6	4.0	1.4	3.0	3.3	6.2	3.7	1.7	4.1	5.5	5.0
	3.0	0.3	0.8	2.8	5.2	3.7	3.2	2.9	4.6	8.9	3.5	3.1	2.8	4.8	6.1
sc 7	5.5	4.3	2.2	2.4	0.2	8.9	1.9	0.3	0.1	0.0	7.5	2.2	0.5	0.2	0.0
	2.0	5.1	9.9	2.0	1.5	7.8	7.0	1.8	2.7	1.0	7.1	9.5	4.5	3.7	0.8
	3.1	0.6	2.4	8.2	9.9	4.8	4.3	4.6	7.6	7.1	4.3	4.5	5.3	6.2	3.5
sc 8	0.8	11.1	4.4	2.7	0.2	5.8	7.1	5.9	0.2	0.0	5.3	7.5	6.2	0.2	0.0
	0.1	0.8	17.1	2.3	1.6	3.4	1.4	14.2	1.0	0.0	3.3	1.6	15.6	1.2	0.0
	3.0	0.1	2.1	10.7	2.3	3.2	3.3	8.1	6.0	0.4	3.1	3.2	6.4	6.0	0.4

In bold are given the MTD(s).

Table 8: Percentage of combination recommendation for multiple MTDs selection using decision rules described in simulation section.

Agent	Agent 1					Agent 1					Agent 1					No MTD																			
	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5
2	Scenario 1																																		
	TSTAT																																		
	3	20.8	49.3	18.3	2.8	0.6	8.2	13.6	61.5	20.6	2.1	0.2	1.8											0.2	1.8										
	2	0.5	15.9	55.5	23.2	4.8	0.0	0.8	15.2	54.9	24.6	4.2	0.1											4.2	0.1										
	1	0.0	3.2	24.1	49.0	21.1	2.5	0.3	6.9	23.9	46.1	20.4	2.2											20.4	2.2										
	BCOPULA																																		
	3	10.7	63.5	17.2	0.9	0.0	7.6	30.8	35.2	18.9	0.5	0.0	14.6											0.0	14.6										
	2	0.2	15.7	71.3	12.2	0.5	0.0	21.5	10.2	47.9	20.1	0.3	0.0											0.3	0.0										
	1	0.0	2.1	29.8	55.6	12.3	0.0	12.2	18.1	24.3	41.0	4.3	0.0											4.3	0.0										
	Scenario 2																																		
TSTAT																																			
3	12.0	1.7	0.0	0.0	0.0	86.2	26.8	4.5	0.4	0.0	0.0	68.2											0.0	68.2											
2	71.2	21.8	4.1	0.4	0.1	2.2	64.8	24.6	7.0	1.4	0.1	2.0											0.1	2.0											
1	11.7	49.6	25.1	8.9	4.5	0.0	18.0	46.4	19.6	9.7	6.2	0.0											6.2	0.0											
BCOPULA																																			
3	43.5 ∞	2.0	0.0	0.0	0.0	54.5	34.5	0.1	0.0	0.0	0.0	65.3											0.0	65.3											
2	84.6	13.8	0.5	0.0	0.0	1.0	82.5	16.6	0.5	0.0	0.0	0.4											0.0	0.4											
1	16.3	76.9	6.4	0.2	0.0	0.0	18.8	70.9	10.3	0.1	0.0	0.0											0.0	0.0											
Scenario 3																																			
TSTAT																																			
3	3.4	21.2	49.5	18.2	2.8	4.8	0.7	14.8	64.4	17.6	1.4	1.0											1.4	1.0											
2	0.0	1.1	16.5	54.1	26.8	1.5	0.0	1.2	17.8	56.8	22.6	1.4											22.6	1.4											
1	0.0	1.7	7.1	34.4	44.4	12.2	0.0	3.3	8.4	30.3	44.6	13.3											44.6	13.3											
BCOPULA																																			
3	1.1	14.1	64.5	16.4	0.4	3.4	31.3	2.9	33.8	16.6	1.9	13.5											1.9	13.5											
2	0.0	0.4	18.9	64.6	16.0	0.0	13.6	2.5	10.8	55.4	17.8	0.0											17.8	0.0											
1	0.0	1.4	2.2	20.9	75.4	0.0	4.8	12.8	10.3	25.1	47.0	0.1											47.0	0.1											
Scenario 4																																			
TSTAT																																			
3	1.1	0.1	0.0	0.0	0.0	98.7	6.1	1.0	0.0	0.0	0.0	92.9											0.0	92.9											
2	43.9	7.9	0.6	0.1	0.0	47.4	46.1	6.8	1.9	0.3	0.0	44.7											0.0	44.7											
1	51.8	33.0	9.3	2.7	0.3	2.8	57.5	26.1	7.4	3.6	1.7	3.5											1.7	3.5											
BCOPULA																																			
3	4.7	0.2	0.0	0.0	0.0	95.0	1.9	0.0	0.0	0.0	0.0	98.1											0.0	98.1											
2	50.0	2.2	0.0	0.0	0.0	47.8	47.4	0.2	0.0	0.0	0.0	52.3											0.0	52.3											
1	90.4	9.5	0.0	0.0	0.0	0.0	73.2	13.4	0.1	0.0	0.0	13.3											0.0	13.3											

In bold are given the MTDs.

Chapter 4

A Bayesian dose-finding design for drug combination clinical trials based on the logistic model

Background: A major part of my PhD dealt with phase I clinical trials in oncology for the combination of molecules. As combination is becoming a standard in oncology, our aim was to develop innovative adaptive designs to answer the current needs in this field. For now, when combining several agents, two cases can be observed: (1) only one agent is varying while the other agents are fixed, or (2) at least two agents are varying. In the first case, the most common in the current practice, the problem is simply brought back to one-dimension, and single-agent designs can be used appropriately. Nevertheless, many combinations are not explored and combinations that could be more appropriate in terms of toxicity and efficacy could be missed. In the second case, true combination designs are nearly never applied and the dose-toxicity relationship is often viewed as a one-dimensional dose space while the reality involved several agents inducing a multi-dimensional issue. To bring the problem back into a one-dimensional space, physicians pre-select the combinations to be evaluated associated with a known toxicity order. In order to reduce the multidimensional combination space into a one-dimensional space, combinations for which the toxicity order is unknown are deleted or an arbitrary ordering is assumed. To explore the entire combination space is obviously not feasible in practice and physicians only wish to explore a subset of combinations. Nevertheless, the choice of the combinations to explore should not be limited by partial toxicity ordering and the design should have the possibility to explore any combination it estimates to be the best. Indeed, due to the possible interactions between drugs, pre-selecting an arbitrary reduced subset of combination induces a risk not to select any combination with a toxicity rate close to the target toxicity.

Methods for single-agent trials are not always appropriate for combination when several agents are varying as they are not designed to take into account the multi-dimensional space. Several alternative designs have been proposed for com-

binations that are either algorithm-based or design-based that give the possibility to explore any appropriate combination in the entire combination space according to the accumulated data. In the context of combination of two cytotoxic agents, in Chapter 3 we reviewed the statistical literature and compared several representative existing methods. Our simulation study highlighted that no design stands out and the performance of the designs are in general comparable with high operational characteristics dependent on the scenarios.

Objective: On this basis, our aim was to propose a dose-finding design that (1) had good operational characteristics in different possible location of the MTD(s), and (2) that in general would perform better than the existing designs and could be therefore recommended for current practice.

Method: A simple and efficient statistical model used in the context of single-agent trials is a logistic model with one or two parameters. Nevertheless, more sophisticated models were used for combination. We decided to model the dose-toxicity relationship of the combination based on a logistic model, that are often well-known by physicians, with both agents and an interaction term. We then focused on the construction of the dose-allocation process and MTD recommendation as these decision rules are of major importance on the performance of dose-finding designs. We used adaptive rejection metropolis sampling (ARMS) within Gibbs sampling (GS) to estimate model parameters (see Appendix A.3). We chose to base the MTD recommendation on intervals on the toxicity probability distribution at each dose level to use not only the estimated mean or median toxicity probability but also the uncertainty around toxicity probabilities estimation.

We have proposed a statistical method for clinical trial designs that evaluate combinations of two agents. Our aim was not to find the correct order of toxicity within the drug combination space but to identify the right dose combination to be evaluated further in terms of efficacy. For the combination of two cytotoxic agent, as both toxicity and efficacy are increasing with the doses, only toxicity need to be studied. In this case, as in single-agent, the objective is to determine the dose-combination closest to the target toxicity.

Results: Our method seems to be able to identify the MTD with a high percentage of correct selections in a wide variety of scenarios. We compared our method with other model-based designs for combination drug trials. All the designs seem to be efficient when the MTDs are located on the same diagonal in the combination space. One benefit of our method compared with the other proposed designs is that it is also efficient when the MTDs are not necessarily located on the same diagonal.

A Bayesian dose-finding design for drug combination clinical trials based on the logistic model

Marie-Karelle Riviere,^{a,b*} Ying Yuan,^c Frédéric Dubois,^b and Sarah Zohar^a

In early phase dose-finding cancer studies, the objective is to determine the maximum tolerated dose, defined as the highest dose with an acceptable dose-limiting toxicity rate. Finding this dose for drug-combination trials is complicated because of drug–drug interactions, and many trial designs have been proposed to address this issue. These designs rely on complicated statistical models that typically are not familiar to clinicians, and are rarely used in practice. The aim of this paper is to propose a Bayesian dose-finding design for drug combination trials based on standard logistic regression. Under the proposed design, we continuously update the posterior estimates of the model parameters to make the decisions of dose assignment and early stopping. Simulation studies show that the proposed design is competitive and outperforms some existing designs. We also extend our design to handle delayed toxicities. Copyright © 2014 John Wiley & Sons, Ltd.

Keywords: Bayesian inference; dose finding; drug combination; oncology; phase I trial

1. INTRODUCTION

For oncologists, the objective of phase I dose-finding studies is to determine the maximum tolerated dose (MTD), defined as the highest dose with a relatively acceptable dose-limiting toxicity (DLT) [1,2]. DLT is usually defined as a toxicity of grade 3 or higher according to the US National Cancer Institute toxicity criteria [3]. In practice, patients included in phase I clinical cancer trials have already been heavily pre-treated, and in many cases, no alternative therapeutic options are available to them. For cytotoxic anti-cancer drugs, a dose-toxicity effect is assumed whereby the higher the dose, the greater the risk of DLT and the greater the efficacy.

Dose finding for drug combination trials is more difficult than that for conventional single-agent trials because of complicated drug–drug interactions. Moreover, when combining several agents, the order of the toxicity probabilities is not fully known. Should investigators wish to gradually increase the acceptable level of toxicity during the trial, the appropriate order in which the doses for the various drugs in the combination should be increased would be of great interest. For instance, when combining two cytotoxic agents, it remains difficult to decide how to escalate or de-escalate the dose combination, even when a partial ordering is known [4,5].

Recently, many phase I dose-finding designs have been proposed for drug combination trials. Thall *et al.* proposed a Bayesian dose-finding method based on a six-parameter model [6]. Wang and Ivanova developed a three-parameter model-based method in which the parameters are estimated using Bayesian inference [7]. Mandrekar *et al.* proposed an approach incorporating the toxicity and efficacy of each agent into the identification of an optimal dosing region for the combination by using a continuation ratio model to separate each agent's toxicity and efficacy curves [8,9]. Yin and Yuan developed a Bayesian adaptive

design based on latent 2×2 tables in which the combination's toxicity probabilities in the two-dimensional space are estimated using a Gumbel-type model [10]. Yin and Yuan extended their method by changing to a copula-type model to simulate the effect of two or more drugs in combination [11]. Bailey *et al.* introduced a second agent as a covariate in a logistic model [12]. Wages *et al.* considered an approach based on the continual reassessment method and taking into account different orderings with partial order between combinations. In this case, the MTD is estimated for the order associated with the highest model-selection criterion [5]. Most of these existing designs rely on complicated statistical models that typically are not familiar to clinicians, which hinder their acceptance and application in practice. In addition, the performance of these designs seems comparable, and there is no consensus which design should be used [13]. As a result, many of the dose-finding clinical trials conducted to evaluate drug combinations still use the conventional '3 + 3' approach, which was developed for single agents and was shown to be inefficient in terms of dose identification [14–18].

In this paper, we propose a Bayesian dose-finding design for drug combination trials based on standard logistic regression. Under the Bayesian paradigm, data monitoring, early stopping, and dose assignment occur continuously throughout the trial by updating the posterior estimates of the model parameters.

^aINSERM, UMRS 1138, Equipe 22, Centre de Recherche des Cordeliers, Université Paris 5, Université Paris 6, Paris, France

^bIRIS (Institut de Recherches Internationales Servier), Suresnes, France

^cDepartment of Biostatistics, The University of Texas MD Anderson Cancer Center, Houston, TX 77030, USA

*Correspondence to: Marie-Karelle Riviere, INSERM, UMRS 1138, Equipe 22, Centre de Recherche des Cordeliers, Université Paris 5, Université Paris 6, Paris, France. E-mail: marie-karelle.riviere@crc.jussieu.fr

Simulation studies show that in general, the proposed design provides better performance than some existing designs. To facilitate the use of the proposed design by clinicians, R package will be developed for implementing the new design.

This manuscript is laid out as follows. In Section 2, we propose the simple statistical method for modeling the toxicity probabilities of the drug combination under evaluation. Moreover, we present the likelihood function and the prior specifications for the unknown parameters. In this section, we also describe our allocation and dose-finding method, as well as propose stopping rules. We conduct extensive simulation studies to examine the operating characteristics of our design in Sections 3 and 4 and conclude with a discussion in Section 5.

2. METHODS

2.1. Statistical method for combination evaluation (LOGISTIC)

2.1.1. Dose-combination model. Let there be a two-drug combination used in a phase I dose-finding clinical trial for which the dose-toxicity relationship is monotonic and increases with the dose levels. Let (j, k) denote the dose level of a combination in which j refers to agent 1 ($j = 1, \dots, J$) and k refers to agent 2 ($k = 1, \dots, K$). Y_i is a Bernoulli random variable, denoting the toxicity that is equal to 1 if DLT occurs in patient i and 0 otherwise ($i = 1, \dots, N$). We assume that $n_{j,k}$ patients are allocated at combination (j, k) and that a total of $t_{j,k}$ DLTs are observed for that combination. We define $\pi_{j,k}$ as the toxicity probabilities of combination (j, k) , θ as the target probability of toxicity, and p_1, \dots, p_J and q_1, \dots, q_K as the respective prior toxicity probabilities of agent 1 and agent 2 taken alone. For simplicity purposes, we refer to the selected combination as the MTD in order to maintain the same designation as in single-agent trials. In this manuscript, we focus on finding one MTD at the end of the trial. We note that in some cases, it is of interest to find multiple MTDs that can be further tested in phase II trials; see Yuan and Yin [19] and Ivanova and Wang [20] for related designs.

Let u_j and v_k denote the 'effective' or standardized doses ascribed to the j th level of agent 1 and k level of agent 2, respectively. We model the drug combination-toxicity relationship using a four-parameter logistic model, as follows:

$$\text{logit}(\pi_{j,k}) = \beta_0 + \beta_1 u_j + \beta_2 v_k + \beta_3 u_j v_k, \quad (1)$$

where β_0 , β_1 , β_2 , and β_3 are unknown parameters that represent the toxicity effect of agent 1 (β_1), that of agent 2 (β_2), and that of the interaction between the two agents (β_3). These parameters are defined such that $\beta_1 > 0$ and $\beta_2 > 0$, ensuring that the toxicity probability is increasing with the increasing dose level of each agent alone, $\forall k, \beta_1 + \beta_3 v_k > 0$ and $\forall j, \beta_2 + \beta_3 u_j > 0$, ensuring that the toxicity probability is increasing with the increasing dose levels of both agents together, and intercept $-\infty < \beta_0 < \infty$. The standardized dose of two agents is defined as $u_j = \log\left(\frac{p_j}{1-p_j}\right)$ and $v_k = \log\left(\frac{q_k}{1-q_k}\right)$, where p_j and q_k are the prior estimates of the toxicity probabilities of the j th dose level of agent 1 and the k th dose level of agent 2, respectively, when they are administered individually as a single agent. Before two agents can be combined, each of them typically has been thoroughly investigated individually. Therefore, there is often rich prior information on p_j 's and q_k 's, and their values can be readily elicited from

physicians. Using the prior information to define standardized dose has been widely used in dose-finding trial designs, and the most well-known example perhaps is the skeleton of the continuous reassessment method (CRM) [21] with a logistic model. Research has shown that this approach improves the estimation stability and trial performance [2,22]. Our definition of u_j and v_k can be viewed as an extension of the skeleton of the CRM (with a logistic model) to drug-combination trials.

2.1.2. Likelihood and posterior inference. Under the proposed model, the likelihood is simply a product of the Bernoulli density, given by

$$L(\beta_0, \beta_1, \beta_2, \beta_3 | \text{data}) \propto \prod_{j=1}^J \prod_{k=1}^K \pi_{j,k}^{t_{j,k}} (1 - \pi_{j,k})^{n_{j,k} - t_{j,k}}.$$

We assume that the prior distributions of the model parameters are independent. For β_0 and β_3 , we assign a vague normal prior $N(0, 10)$ centered at 0 to indicate that *a priori* we do not favor either positive or negative values for these parameters and let the observed data speak for themselves through posteriors. For β_1 and β_2 , we assume an informative prior distribution $\text{Exp}(1)$ centered at 1, as those parameters should not be too far from 1. Then, the joint posterior distribution of parameters $\beta_0, \beta_1, \beta_2$, and β_3 is given by

$$f(\beta_0, \beta_1, \beta_2, \beta_3 | \text{data}) \propto L(\beta_0, \beta_1, \beta_2, \beta_3 | \text{data}) f(\beta_0) f(\beta_1) f(\beta_2) f(\beta_3). \quad (2)$$

We sample this posterior distribution using Gibbs sampler, which sequentially draws each of the parameters from their full conditional distributions (see Appendix, available online as Supporting Information). These full conditional distributions do not have closed forms, and we use the adaptive rejection Metropolis sampling (ARMS) method [23] to sample them. In order to impose the constraint that $\forall k, \beta_1 + \beta_3 v_k > 0$ and $\forall j, \beta_2 + \beta_3 u_j > 0$, at each iteration of Gibbs sampling, if sampled β_1, β_2 , and β_3 fail to satisfy the constraint, we re-sample β_1, β_2 , and β_3 . Let $(\beta_0^{(\ell)}, \beta_1^{(\ell)}, \beta_2^{(\ell)}, \beta_3^{(\ell)})_{\ell=1, \dots, L}$ denote the L posterior samples obtained from Gibbs sampler; the posterior toxicity probabilities can be estimated using Monte Carlo by

$$\tilde{\pi}_{j,k} = \frac{1}{L} \sum_{\ell=1}^L \frac{\exp(\beta_0^{(\ell)} + \beta_1^{(\ell)} u_j + \beta_2^{(\ell)} v_k + \beta_3^{(\ell)} u_j v_k)}{1 + \exp(\beta_0^{(\ell)} + \beta_1^{(\ell)} u_j + \beta_2^{(\ell)} v_k + \beta_3^{(\ell)} u_j v_k)}.$$

2.1.3. Dose-finding algorithm and determination of the MTD. During the trial conduct, we use the dose-finding algorithm proposed by Yin and Yuan [10,11] to determine dose escalation and de-escalation, and propose a different criterion for MTD selection at the end of the trial. In our design, similar to Yin and Yuan [11], we restrict dose escalation and de-escalation one level at a time (i.e., we do not allow dose escalate or de-escalate along the diagonal) based on the practical consideration that physicians are conservative and typically do not allow two agents to escalate at the same time for patient safety. Nevertheless, we note that Sweeting and Mander [24] showed that diagonal escalation strategy may be more efficient in reaching the target toxicity level quicker with fewer patients treated at sub-optimal doses and have a higher percentage of correct selection at the end of the trial. We take this strategy in our start-up phase described later.

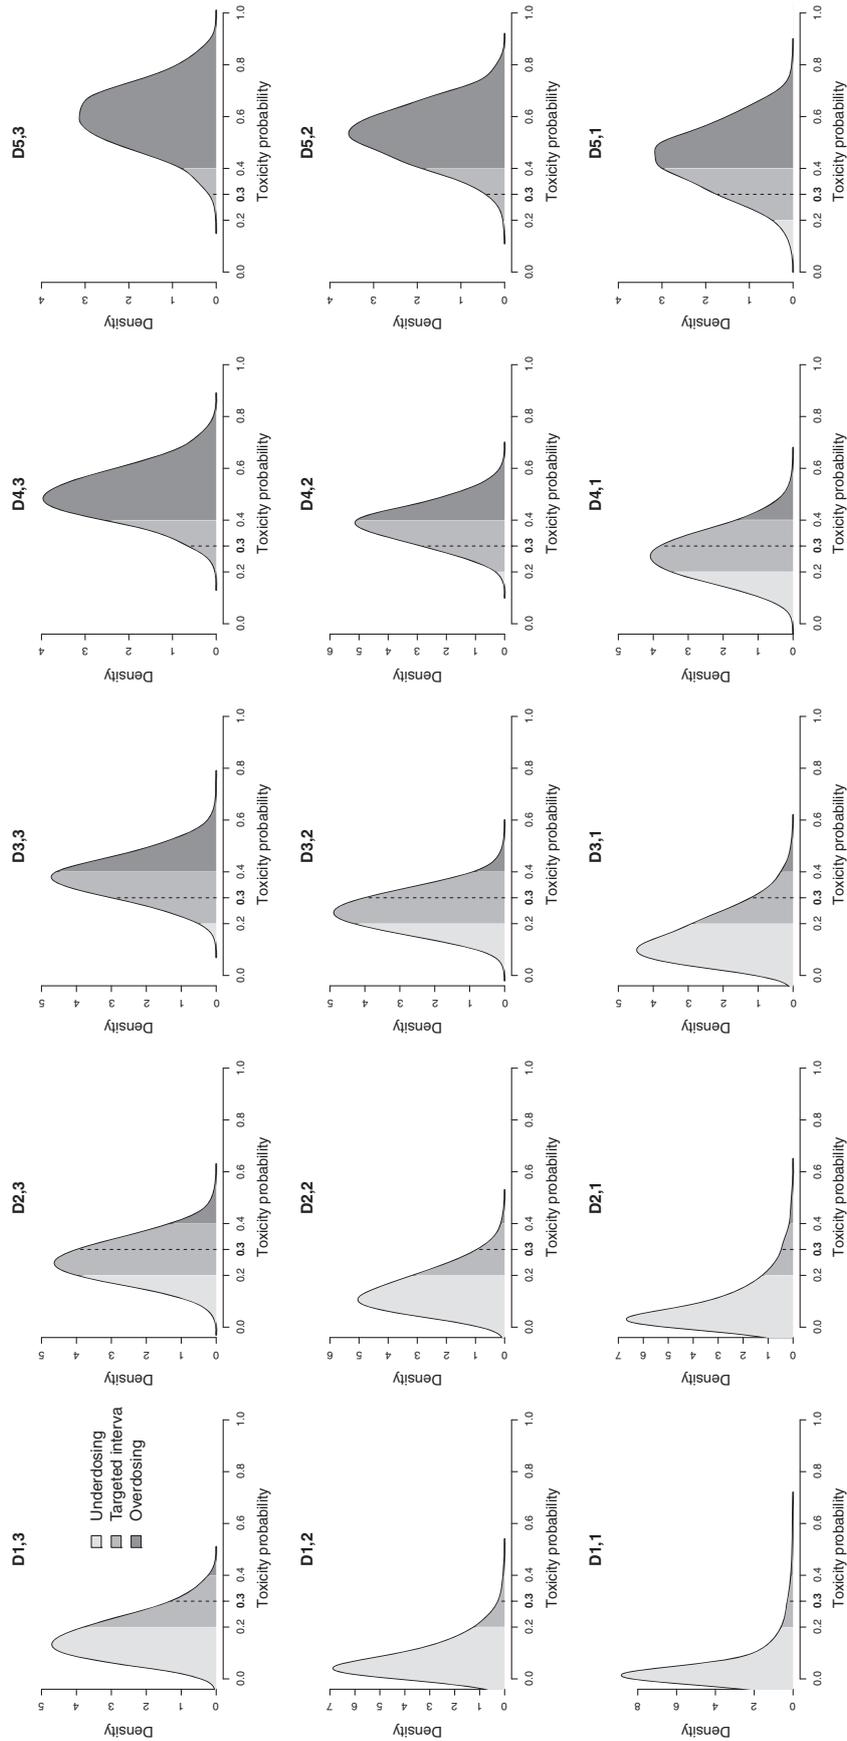


Figure 1. Posterior densities of the toxicity probability for each combination. The dashed line represents the toxicity target.

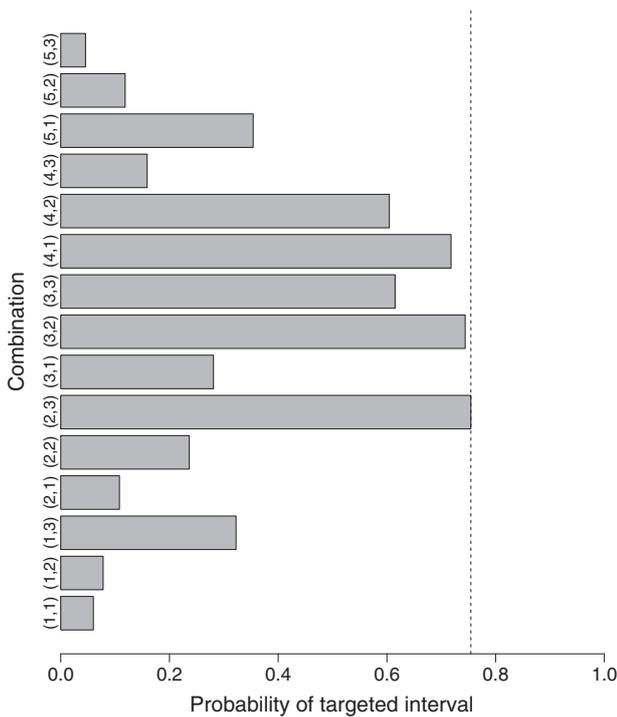


Figure 2. Probability of achieving the targeted interval, $P(\pi_{j,k} \in [\theta - \delta; \theta + \delta])$ at the end of the trial for each combination. The dashed vertical line corresponds to the highest probability, which determines the combination selected as the MTD ((2, 3) in this example).

Let c_e be the probability threshold for dose escalation and c_d the probability threshold for dose de-escalation. We require $c_e + c_d > 1$ to avoid that the decisions of dose escalation and deescalation occur at the same time. Our dose-finding algorithm can be described as follows:

- If the current combination is (j, k) and $P(\pi_{j,k} < \theta | \text{data}) > c_e$, we escalate the combination dose level to an adjacent combination dose level $\{(j + 1, k), (j, k + 1), (j + 1, k - 1), (j - 1, k + 1)\}$ that has a toxicity probability that is higher than the current value and closest to θ . If the current combination is the highest of the combination space, (J, K) , we retain the same combination dose level for the next cohort.
- If the current dose combination is (j, k) and $P(\pi_{j,k} > \theta | \text{data}) > c_d$, we de-escalate the combination dose level to an adjacent combination dose level $\{(j - 1, k), (j, k - 1), (j + 1, k - 1), (j - 1, k + 1)\}$ that has a toxicity probability that is lower than the current value and closest to θ . If the current combination is the lowest of the combination space, $(1, 1)$, we retain the same combination dose level for the next cohort.
- If the current combination is (j, k) and $P(\pi_{j,k} < \theta | \text{data}) \leq c_e$ and $P(\pi_{j,k} > \theta | \text{data}) \leq c_d$, we treat the next cohort of patients at the current combination dose level.

Once we reach the maximum sample size, we select the MTD as the combination associated with the highest posterior probability, $P(\pi_{j,k} \in [\theta - \delta; \theta + \delta])$, and which have been used to treat at least one cohort of patients. If, for example, the target toxicity θ is 0.3 and the length around the targeted interval is $\delta = 0.1$, which gives a targeted interval defined as $[0.2; 0.4]$, then, as illustrated in Figure 1, at the end of the trial, we can obtain the posterior

densities of the toxicity probability for each combination. Shown in light gray in the figure is the area under the curve (AUC) for a toxicity probability lower than 0.2, which is equal to the probability of under-dosing. Shown in medium gray is the AUC for a toxicity probability between 0.2 and 0.4, which is the probability of being in the targeted interval. Shown in dark gray is the AUC for a toxicity probability greater than 0.4, which is equal to the probability of overdosing. For each combination, given the probabilities of being in the targeted interval, the combination already administered to at least one cohort at the end of the trial and corresponding to the highest probability is selected as the MTD, for example, dose combination (2, 3), as illustrated in Figure 2.

The probability thresholds c_e and c_d are critical for the performance of the design as they control the dose escalation and de-escalation. The values of c_e and c_d should be carefully calibrated through simulation to ensure good operating characteristics of the design. In practice, this can be performed as follows: First, define a set of representative dose-toxicity scenarios that may be encountered in the trial, and then conduct simulation under different values of c_e and c_d to evaluate the performance of the design. This is a trial-and-error process and may involve repeatedly tuning the values of c_e and c_d based on the simulation results. The goal is to find the values of c_e and c_d that yield good overall performance across different scenarios (e.g., the percentage of correct selection of the MTD, the number of patients exposed to over-toxic combinations or under-toxic combinations). Such a calibration-based approach has been widely used in clinical trial designs [6,10,11,25].

Because of the limited availability of information at the beginning of the trial, the posterior estimates of the toxicity probabilities may not be reliable. Therefore, as suggested by other authors [10,16,20,26,27], we implement an algorithm-based start-up phase in order to gather enough information to estimate the $\pi_{j,k}$. Our start-up phase shares the spirit of accelerated titration design [28] and can be described as follows: Treat the first cohort of patients at the lowest dose combination (1, 1); if no toxicity is observed, escalate the dose along the diagonal, and treat the second cohort of patients at (2, 2); we continue this dose escalation along the diagonal until the observation of a toxicity or at least one agent is at its maximum. If one of the agents is already at its maximum dose level and still no toxicity is observed, we increase the dose of the other agent until both agents reach their maximum dose level. At any time of the start-phase, once the first toxicity event is observed, the start-up phase is completed, and the design switches to the model-based dose-finding algorithm as described earlier.

2.1.4. Stopping rules. Our aim is to propose a design that, should the investigators wish, will stop the trial when all combinations are estimated to be unacceptable in terms of toxicity. If the current combination is (1, 1), at least two cohorts have been included and $P(\pi_{j,k} < \theta) \geq 0.975$, then the trial is stopped.

2.2. Some existing designs for combination studies

2.2.1. BCOPULA and BGUMBEL methods. Yin and Yuan proposed two Bayesian methods that use copula regression for combinations. The first method [11] uses a Clayton-copula regression in

order to express the joint probability of combination (j, k) with the marginal true probabilities of toxicity (p_j^α, q_k^β) :

$$\pi_{j,k} = 1 - \left((1 - p_j^\alpha)^{-\gamma} + (1 - q_k^\beta)^{-\gamma} - 1 \right)^{-\frac{1}{\gamma}}.$$

The second design [10] uses a Gumbel model:

$$\pi_{j,k} = 1 - (1 - p_j^\alpha) (1 - q_k^\beta) \left[1 + p_j^\alpha q_k^\beta \frac{e^\gamma - 1}{e^\gamma + 1} \right],$$

where $\gamma, \alpha, \beta > 0$ are unknown parameters. The parameter γ characterizes the drug interactive effect, and α and β characterize the uncertainty of the initial guesses. The combination allocation algorithm is the same as that presented in Section 2.1. The final MTD is the combination with a toxicity probability closest to the target among the combinations already administered in the trial.

2.2.2. I2D method. Wang and Ivanova proposed the I2D method [7], which is a two-dimensional model-based design, defined as follows [7]:

$$\pi_{j,k} = 1 - (1 - p_j)^\alpha (1 - q_k)^\beta \exp(-\gamma \log(1 - p_j) \log(1 - q_k)),$$

where $\alpha > 0, \beta > 0$, the interaction term $\gamma > 0$, and p_j, q_k are the working models for agent 1 and agent 2, respectively.

The combination that will be allocated to the next cohort is the combination closest to the target that belongs to $\{(j + 1, k), (j, k + 1), (j - 1, k + 1), (j - 1, k), (j, k - 1), (j + 1, k - 1), (j, k)\}$. The final MTD is defined as the combination with a toxicity probability closest to the target among the combinations already administered in the trial.

2.2.3. LOGODDS. Gasparini *et al.* [29] have proposed an alternative model, $\pi_{j,k}$, to the copula proposed by Yin and Yuan [11] containing an interaction parameter γ that explicitly quantifies departure from no interaction on the log-odds scale:

$$\gamma = \text{logodds}(\pi_{j,k}) - \text{logodds}(\pi_{j,k}^\perp)$$

where $\pi_{j,k}^\perp$ is the no-interaction model defined as $\pi_{j,k}^\perp = p_j + q_k - p_j q_k$. By back-solving this equation, an explicit expression for the probability of toxicity is obtained. A normal prior centered in 0 and with variance of 100 was chosen for the interaction parameter. The rest of the dose allocation process, estimation, and MTD determination was the same as our proposed design in order to compare our method involving a simple interaction logistic model with other logistic models.

2.2.4. TMML. Thall *et al.* [6] proposed a six-parameter model defined as follows:

$$\pi_{j,k} = \frac{\alpha_1 x_1^{\beta_1} + \alpha_2 x_2^{\beta_2} + \alpha_3 (x_1^{\beta_1} x_2^{\beta_2})^{\beta_3}}{1 + \alpha_1 x_1^{\beta_1} + \alpha_2 x_2^{\beta_2} + \alpha_3 (x_1^{\beta_1} x_2^{\beta_2})^{\beta_3}}$$

where $\alpha_1, \beta_1, \alpha_2, \beta_2, \alpha_3, \beta_3$ are unknown positive parameters. Following Thall *et al.* [6], a prior gamma distribution with parameters (0.3, 0.3) was chosen for α_1, α_2 , and α_3 , and a gamma distribution with parameters (0.003, 0.03) was chosen for β_1, β_2 ,

and β_3 . The rest of the dose allocation process, estimation, and MTD determination was the same as our proposed design in order to compare our method involving a simple interaction logistic model with other logistic models.

2.2.5. One-dimensional CRM. In practice, one-dimensional CRM sometimes is used to conduct dose-combination trials [13]. Under this method, we first preselect a subset of combinations, for which the toxicity probability order is known, and then apply the standard CRM to find the MTD. The drawback of such an approach is that we only investigate a subset of the whole two-dimensional dose space and may miss the true MTD. Following [13], we chose the subset of combinations as the combinations located at anti-diagonal path: $D_{1,1} \rightarrow D_{1,2} \rightarrow D_{2,2} \rightarrow D_{3,2} \rightarrow D_{4,2} \rightarrow D_{5,2} \rightarrow D_{5,3}$.

3. SIMULATIONS

We simulated 2000 independent replications of phase I trials that evaluate two agents in drug combinations, with five dose levels for agent 1 and three for agent 2, giving 15 possible combinations. We simulated 14 scenarios to represent the possible true underlying combination toxicities (Table I). We studied several locations of the MTD in the combination space, as well as a number of correct combinations (or MTDs) that could shift from 0 to 3. We fixed the toxicity target at 0.3 and used an overall sample size of 60. In this paper, we compare the performance of our method with those of five other designs, including the BCOPULA [11] and BGUMBEL [10] methods proposed by Yin and Yuan, and the I2D method proposed by Wang and Ivanova [7], and with two other logistic models, LOGODDS proposed by Gasparini *et al.* [29] and TMML proposed by Thall *et al.* [6]. We also compared the performance of this multidimensional designs with a one-dimensional CRM, CRM anti-diag, as described previously. The design parameters of all the designs (e.g., working model) have been calibrated via simulation before used for the comparison.

To ensure comparability, we chose a cohort size of 3 for all methods and did not use stopping rules when comparing all the methods. The start-up phase was implemented in all simulations. Because of ethical concerns, we started each trial at the lowest dose combination (1, 1). We selected these features in order to employ typical trial set-ups in our simulation study. For agent 1, we specified the marginal initial guesses of toxicities, p_j , as (0.12, 0.2, 0.3, 0.4, 0.5), and for agent 2, we specified q_k as (0.2, 0.3, 0.4). The same working models were used for BCOPULA, BGUMBEL, LOGODDS, and I2D. We set the length around the targeted interval, δ , at 0.1. Based on a sensitivity analysis (data not shown), we set the probability thresholds c_e and c_d at 0.85 and 0.45. We recorded 5000 posterior samples of the model parameters after 2000 burn-in iterations to make inference. After comparing the proposed design to the other designs, we also investigated the performance of our design with the stopping rule that we introduced earlier. For the one-dimensional CRM, we used the 'dfcrm' R package with restrictions to avoid skipping doses. The working model was generated using the 'getprior' function with an indifference interval $\delta = 0.05$, an initial guessed MTD at dose level 4 for a trial with seven doses, and the same settings as for multidimensional designs were used.

Table I. Toxicity scenarios for the two-agent combinations.

Agent 2	Agent 1									
	1	2	3	4	5	1	2	3	4	5
Scenario 1										
3	0.15	0.30	0.45	0.50	0.60	0.45	0.55	0.60	0.70	0.80
2	0.10	0.15	0.30	0.45	0.55	0.30	0.45	0.50	0.60	0.75
1	0.05	0.10	0.15	0.30	0.45	0.15	0.30	0.45	0.50	0.60
Scenario 2										
3	0.10	0.15	0.30	0.45	0.55	0.50	0.60	0.70	0.80	0.90
2	0.07	0.10	0.15	0.30	0.45	0.45	0.55	0.65	0.75	0.85
1	0.02	0.07	0.10	0.15	0.30	0.30	0.45	0.60	0.70	0.80
Scenario 3										
3	0.07	0.09	0.12	0.15	0.30	0.15	0.30	0.45	0.50	0.60
2	0.03	0.05	0.10	0.13	0.15	0.09	0.12	0.15	0.30	0.45
1	0.01	0.02	0.08	0.10	0.11	0.05	0.08	0.10	0.13	0.15
Scenario 4										
3	0.30	0.50	0.60	0.65	0.75	0.08	0.15	0.45	0.60	0.80
2	0.15	0.30	0.45	0.52	0.60	0.05	0.12	0.30	0.55	0.70
1	0.07	0.10	0.12	0.15	0.30	0.02	0.10	0.15	0.50	0.60
Scenario 5										
3	0.30	0.37	0.42	0.47	0.52	0.08	0.10	0.15	0.30	0.50
2	0.15	0.30	0.37	0.43	0.48	0.04	0.07	0.12	0.16	0.30
1	0.10	0.12	0.30	0.40	0.45	0.01	0.03	0.06	0.08	0.10
Scenario 6										
3	0.50	0.60	0.70	0.80	0.90	0.30	0.42	0.52	0.62	0.70
2	0.10	0.30	0.50	0.70	0.80	0.10	0.20	0.30	0.40	0.50
1	0.06	0.10	0.15	0.30	0.50	0.05	0.12	0.20	0.30	0.40
Scenario 7										
3	0.42	0.52	0.62	0.70	0.80	0.30	0.42	0.52	0.70	0.80
2	0.20	0.30	0.40	0.50	0.67	0.10	0.20	0.30	0.50	0.67
1	0.12	0.20	0.30	0.40	0.60	0.04	0.06	0.08	0.20	0.30
Scenario 8										
Scenario 9										
Scenario 10										
Scenario 11										
Scenario 12										
Scenario 13										
Scenario 14										

The true MTD combinations are shown in boldface.

4. RESULTS

For each scenario, we present the correct MTD selection rate, or percentages of correct selection (PCS), in Table II. In general, the proposed method for combination evaluation (LOGISTIC) performed better than the other existing model-based designs. Indeed, except in scenarios 4 and 5, in which the PCS were already very high (86.7% and 80.4%, respectively), those in other scenarios were either the best or not less than 2% of those of the other designs. Moreover, the PCS were greater than 55% in all scenarios and higher than 60% in 12 scenarios out of 14.

Scenarios 1, 3, and 9 included three possible MTDs that were located on one diagonal of the combination space: combinations (2, 3), (3, 2), (4, 1) in scenario 1, combinations (3, 3), (4, 2), (5, 1) in scenario 3, and combinations (1, 3), (2, 2), (3, 1) in scenario 9 (Table I). For these scenarios, all model-based designs gave high PCS, whereas LOGISTIC seemed to perform better than the other methods. In addition, for the LOGISTIC, the distribution of allocated dose levels around the MTDs (Table III) seemed to be spread unevenly between possible MTDs.

In scenarios 2 and 10, in which two possible MTDs were located on the same diagonal at the lower end or at the higher end of the combination space, high PCS were observed for the LOGISTIC (> 75%). When the correct combination was in the lower (1, 1) extremity of the combination space, as in scenario 4, the TMML model and then I2D method performed best (92.1% and 89.7%, respectively), but the PCS for the LOGISTIC were close to this value (86.7%). By contrast, when the correct combination was in the higher (5, 3) extremity of the combination space, as in scenario 5, the LOGODDS model gave the highest PCS with 90.0%.

In scenarios 6, 7, 11, 12, 13, and 14, when the true MTDs were not located on the same diagonal but randomly located in the combination space, the proposed method performed best with the PCS that was consistently higher than 56%. In contrast, the PCS of other methods can be as low as 30%. For example, under scenario 6, the PCS of BCOPULA, BGUMBEL, I2D, and LOGODDS were all lower than 47%, and under scenario 12, the PCS of TMML was only 47.9%.

Table II. Percentage of correct selection in comparison of six dose-finding designs. Percentage of patients allocated to a true MTD during the trial and mean number of observed DLTs throughout the trial.														
Scenario	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8	sc9	sc10	sc11	sc12	sc13	sc14
	Percentages of correct selection (PCS)													
CRM anti-diag	73.7	74.8	71.9	84.9	80.0	71.4	73.2	84.3	58.4	74.8	82.8	61.0	59.4	73.2
BCOPULA	66.2	71.8	71.7	84.1	78.1	30.7	49.6	43.5	66.5	60.9	42.7	57.3	44.8	42.3
BGUMBEL	67.1	72.5	68.4	87.5	77.9	33.6	48.0	49.5	64.7	66.2	51.2	56.3	48.7	42.9
I2D	68.0	73.7	66.9	89.7	83.7	37.2	41.9	50.4	63.8	66.7	47.1	45.3	41.6	47.4
LOGODDS	62.8	72.7	63.7	87.5	90.0	46.7	73.2	12.1	57.3	68.4	61.1	56.5	47.0	54.5
TMMIL	70.1	70.5	65.8	92.1	83.9	49.6	65.6	69.7	61.7	66.0	75.9	47.9	60.1	56.0
LOGISTIC	75.4	80.5	74.9	86.7	80.4	63.7	71.2	56.9	69.6	75.1	77.8	56.7	60.0	61.0
	Percentage of patients allocated to a true MTD during the trial													
CRM anti-diag	49.2	55.4	45.2	74.5	45.5	43.5	52.6	58.4	44.1	42.6	59.8	42.1	43.6	49.3
BCOPULA	40.0	50.1	40.3	84.1	27.8	16.6	38.3	23.6	46.2	32.2	34.3	37.5	26.7	33.8
BGUMBEL	40.8	52.8	39.5	81.6	30.5	20.0	34.1	26.0	46.2	34.3	34.2	37.3	31.5	32.0
I2D	44.1	55.6	38.9	79.8	34.6	23.0	32.0	24.0	46.2	30.8	33.5	32.2	29.9	32.4
LOGODDS	40.2	45.3	40.3	82.0	62.5	29.2	47.8	9.8	38.2	43.2	42.8	25.2	31.3	24.8
TMMIL	34.5	40.3	33.7	89.8	42.0	22.2	33.5	24.8	36.5	30.0	35.3	25.0	27.0	26.7
LOGISTIC	44.0	50.5	46.0	78.0	40.0	33.0	43.5	24.3	47.3	38.0	44.0	33.0	40.7	34.3
	Mean number of observed DLTs throughout the trial													
CRM anti-diag	16.5	18.4	15.3	20.3	11.5	14.8	17.6	17.1	17.0	14.1	18.0	16.2	17.0	16.5
BCOPULA	14.2	16.1	12.7	19.5	9.2	12.8	14.6	14.3	15.0	11.5	15.5	14.5	15.2	14.5
BGUMBEL	14.6	16.6	13.2	19.7	9.5	13.5	15.6	14.7	15.4	11.9	16.1	15.0	15.5	15.2
I2D	15.3	17.6	14.1	19.9	10.1	14.3	16.1	15.3	16.2	12.9	16.6	15.6	16.6	15.8
LOGODDS	16.8	17.5	16.4	20.0	13.7	16.6	17.0	16.3	17.1	16.0	16.7	16.8	16.8	16.8
TMMIL	13.8	15.7	13.0	19.2	11.4	13.1	14.2	13.9	14.5	12.4	14.7	14.1	14.8	13.9
LOGISTIC	15.2	17.7	14.1	20.4	11.4	14.3	15.5	15.3	15.9	12.7	16.2	15.7	16.4	15.4

In the top third of the table, the best PCS and those associated with 2% less are shown in boldface.
CRM, continuous reassessment method; MTD, maximum tolerated dose; DLTs, dose-limiting toxicities.

Table III. Combination selection percentages of the proposed method for combination evaluation.

Agent 2	Agent 1									
	1	2	3	4	5	1	2	3	4	5
Scenario 1										
3	6.2	27.5	7.1	0.1	0.0	0.2	0.1	0.0	0.0	0.0
2	0.1	4.6	41.7	2.9	0.1	35.4	11.2	0.1	0.0	0.0
1	0.0	0.1	3.2	6.2	0.1	5.8	45.1	2.1	0.0	0.0
Scenario 2										
Scenario 3										
3	0.5	6.2	43.1	7.5	0.1	0.0	0.1	0.0	0.0	0.0
2	0.0	0.2	4.5	29.1	2.0	4.2	0.1	0.0	0.0	0.0
1	0.0	0.0	0.1	4.0	2.7	86.7	9.0	0.0	0.0	0.0
Scenario 4										
Scenario 5										
3	0.0	0.1	1.5	8.5	80.4	3.0	44.0	14.1	1.4	0.1
2	0.0	0.0	0.1	1.0	7.8	0.0	1.2	11.1	19.7	2.4
1	0.0	0.0	0.0	0.1	0.5	0.0	0.0	0.2	2.1	0.5
Scenario 6										
Scenario 7										
3	15.8	2.4	0.1	0.0	0.0	0.9	15.5	17.0	0.0	0.0
2	4.2	54.4	15.6	1.1	0.0	0.0	0.9	56.9	0.5	0.0
1	0.0	0.5	3.4	1.6	1.0	0.0	0.0	2.6	5.7	0.0
Scenario 8										
Scenario 9										
3	8.1	5.7	1.8	0.2	0.0	0.1	0.4	7.1	43.4	8.2
2	6.4	48.2	11.7	1.5	0.0	0.0	0.0	0.3	8.1	31.7
1	0.1	1.8	13.3	1.4	0.0	0.0	0.0	0.0	0.1	0.8
Scenario 10										
Scenario 11										
3	6.7	0.2	0.0	0.0	0.0	22.7	12.5	0.5	0.0	0.0
2	2.4	75.3	6.0	0.1	0.0	0.4	21.3	31.4	2.1	0.0
1	0.0	0.2	6.6	2.5	0.1	0.0	0.2	6.2	2.6	0.0
Scenario 12										
Scenario 13										
3	7.0	0.8	0.0	0.0	0.0	24.5	12.7	0.6	0.0	0.0
2	13.7	48.0	5.3	0.0	0.0	0.1	18.1	35.6	3.3	0.1
1	0.4	11.9	12.0	0.9	0.0	0.0	0.0	0.7	3.5	0.9
Scenario 14										

The percentages for the true MTD combinations are shown in boldface.

Table IV. Additional toxicity scenario for the two-agent combinations.

Agent 2	Agent 1				
	1	2	3	4	5
Scenario 15					
3	0.60	0.65	0.75	0.85	0.90
2	0.55	0.60	0.70	0.75	0.80
1	0.45	0.50	0.55	0.60	0.65

Finally, when the true MTD was unique and located in the middle of the combination space, as in scenario 8, the PCS decreased. The TMML gave the best PCS (69.7%); nevertheless, the LOGISTIC was the second most efficient design and gave good results with PCS of 56.9% for that scenario (Table II). The PCS of the LOGODDS was merely 12.1%.

Under our simulation setting, the one-dimensional CRM (based on an anti-diagonal subset of combinations) yielded higher PCS than the multidimensional designs. However, this result should not be generalized. The reason that the one-dimensional CRM performed very well is simply because the pre-selected subset of combinations happened to include the true MTD. As the CRM focused on only a subset of doses, it had more patients per dose to find the target doses than the multidimensional designs, given the same total number of patients. It is easy to see that if the true MTD(s) is (are) not contained in the pre-selected subset of ordered combinations, then the one-dimensional CRM will perform badly and can never select the true MTD.

The mean number of DLTs observed throughout the trial in each scenario when using the LOGISTIC was similar to that observed when using the BCOPULA, BGUMBEL, and I2D methods (Table II). Over all the scenarios, the mean number of DLTs observed during the trial was 15.4, which is lower than 18, the expected number of DLTs for 60 patients with a 0.30 toxicity probability.

Table V. Percentage of correct selection (PCS) and percentage of stopping for toxicity without MTD recommendation. Results of the sensitivity analysis with different prior distributions, using for β_0, β_3 normal distribution with high or low variability and for β_1, β_2 exponential or gamma priors with high or low variability.

Scenario	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8	sc9	sc10	sc11	sc12	sc13	sc14	sc15
$\beta_0, \beta_3 \sim N(0, 10), \beta_1, \beta_2 \sim \Gamma(1)$															
PCS	76.7	79.5	74.9	69.8	80.4	64.6	70.6	57.6	68.3	75.1	77.8	56.7	60.7	61.0	–
Stop	0.0	0.7	0.0	16.8	0.0	0.0	0.1	0.0	0.1	0.0	0.0	0.0	0.4	0.0	83.7
$\beta_0, \beta_3 \sim N(0, 50), \beta_1, \beta_2 \sim \Gamma(0.1, 0.1)$															
PCS	72.3	81.8	71.0	68.9	74.8	63.4	70.4	57.9	70.4	72.2	75.1	57.0	58.0	60.9	–
Stop	0.0	0.9	0.0	17.8	0.0	0.0	0.1	0.0	0.1	0.0	0.0	0.0	0.4	0.0	83.5
$\beta_0 \sim N(0, 50), \beta_1, \beta_2 \sim \Gamma(0.5, 0.5), \beta_3 \sim N(0, 10)$															
PCS	75.0	82.1	70.4	70.0	79.0	62.0	70.0	60.3	68.6	72.8	79.4	55.0	57.5	61.6	–
Stop	0.0	0.9	0.0	17.6	0.0	0.0	0.1	0.0	0.1	0.0	0.0	0.0	0.5	0.0	84.1

According to Table II, the LOGISTIC tended to treat as many patients at the MTD combination as the other methods. The lowest percentage of patients treated at the MTD combination was 24.3% in scenario 8, and the highest was 78.0% in scenario 4. In general, the true MTD combination was assigned to at least one-third of the patients, and the mean overall percentage of patients treated at the MTD combination was 42.6%.

After employing the stopping rule for all scenarios when using the LOGISTIC, we added a scenario in which all the dose combinations were unacceptably toxic (Table IV). The addition of the stopping rule for unacceptable toxicity resulted in PCS that were similar to those presented earlier (Table V), except in scenario 4 in which the first dose combination, (1, 1), was the MTD. For this scenario, the PCS decreased by 16.9%, although it was nevertheless high (69.8%). In practice, the trial would not be completely stopped under these circumstances, but lower dose combinations would be added to the trial. In the additional scenario 15, where all dose combinations were too toxic, the trial was stopped in 83.7% of the cases. Therefore, the stopping rule seemed efficient.

4.1. Sensitivity analysis

We conducted a sensitivity analysis in order to study the performance of our design using different prior distributions and parameters values. We used normal distributions centered at 0 for β_0, β_3 with high or low variance (from 10 to 50), and exponential or gamma distributions centered at 1 for β_1, β_2 with high or low variance (from 1 to 10). According to Table V, we can see that the PCS for all scenarios were very similar under these different prior distributions.

4.2. Time-to-event outcome

In practice, a longer follow-up time may be required to assess the toxicity outcome. Therefore, the toxicity outcomes of some patients already treated in the trial may be unobserved (or censored) when a new patient is enrolled in the trial and is ready for dose assignment. Waiting to assess toxicity outcomes for each cohort before including a new one in the trial can greatly increase the duration of the trial. To overcome this issue, we extended our method by modeling toxicity as a time-to-event outcome. Following the Time-to-event Continual Reassessment Method (TITE-CRM) [30], we considered a weighted dose-toxicity relationship $\hat{w}F(D, \beta_0, \beta_1, \beta_2, \beta_3)$ where \hat{w} is monotone and increasing with patient follow-up time such that $0 \leq \hat{w} \leq 1$, and the toxicity probability model $F(D, \beta_0, \beta_1, \beta_2, \beta_3) = \pi_{j,k}$ is the same as that in (1).

Let T be a (maximum) fixed time window during which patients are followed, and $y_{i,N}$, $C_{i,N}$, and $\hat{w}_{i,N}$, respectively, be the indicator of a DLT for patient i prior to the entry of the $(N + 1)^{\text{th}}$ patient, the follow-up time of patient i prior to the entry of the $(N + 1)^{\text{th}}$ patient, and the weight assigned to patient i prior to the entry of the $(N + 1)^{\text{th}}$ patient. Let X_i denote the time to toxicity of the i^{th} patient, and (j_i, k_i) the combination received by patient i . Before combination assignment, the likelihood is defined as

$$L(\beta_0, \beta_1, \beta_2, \beta_3 | \text{data}) = \prod_{i=1}^N (\hat{w}_{i,N} \pi_{j_i, k_i})^{y_{i,N}} (1 - \hat{w}_{i,N} \pi_{j_i, k_i})^{1 - y_{i,N}}$$

Table VI. Percentage of correct selection (PCS) and percentage of stopping for toxicity without MTD recommendation when toxicity is considered as a time-to-event outcome.

Scenario	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8	sc9	sc10	sc11	sc12	sc13	sc14	sc15
TITE-LOGISTIC															
PCS	74.0	77.5	74.1	68.0	78.2	63.3	66.5	57.6	64.7	74.6	76.2	56.5	60.4	62.7	-
Stop	0.1	0.8	0.0	16.6	0.0	0.1	0.1	0.0	0.2	0.0	0.1	0.1	0.2	0.0	82.4

Following Cheung and Thall [31], we chose the weights $\hat{w}_{i,N}$ as

$$\hat{w}_{i,N} = \frac{\#\{m/X_m \leq C_{i,N}, C_{m,N} \geq T\} + \hat{w}_{i,N}^0}{\#\{m/X_m \leq T, C_{m,N} \geq T\} + 1},$$

where m refers to patient, $\#\{m/X_m \leq T, C_{m,N} \geq T\}$ is the number of completely followed patients on who a toxicity was observed, and $\hat{w}_{i,N}^0$ is the linear weight for patient i defined as $\hat{w}_{i,N}^0 = C_{i,N}/T$, that is, the proportion of time patient i was followed compared with the full follow-up time.

We simulated the time-to-toxicity outcomes using an exponential distribution such that the toxicity probabilities at the end of follow-up matched those given in Table I. Using the same working model, we chose a full follow-up time of 3 weeks and simulated patient accrual using a Poisson process with parameter 1, meaning that an average of one patient arrived every week.

Table VI shows the results of the extended LOGISTIC for all 15 scenarios. We observe that the performance of the design decreases only slightly by 2%, and the PCS for all scenarios are still very high. This demonstrates that the LOGISTIC can be used in trials when toxicity cannot be assessed quickly.

5. DISCUSSION

We have proposed a statistical method for clinical trial designs that evaluate various dose combinations for two agents. This method seems to be capable of identifying the MTD with a high percentage in all scenarios. Indeed, this method works for a wide variety of drug combinations. Our aim was not to find the correct order of toxicity within the drug combination space but to identify the right dose combination to be evaluated further in terms of efficacy. We compared our method with five other model-based designs for combination drug trials. All the designs seem to be efficient when the MTDs are located on the same diagonal in the combination space. One benefit of our method compared with the other proposed designs is that it is also efficient when the MTDs are not necessarily located on the same diagonal.

We have designed our method to enable the selection of only one MTD at the end of the trial. There are various reasons why investigators may wish to select one MTD per level of agent 2. For this purpose, our method can be modified by selecting, for each level k of agent 2, the combination (j_k, k) such that $j_k = \text{argmax}_{1 \leq j \leq J} P(\pi_{j,k} \in [\theta - \delta; \theta + \delta])$. In the case where it cannot be assumed that one MTD exists per level of agent 2, a minimum probability of being in the targeted toxicity interval or a maximum probability of overdosing can be added in the determination of the MTD so as not to recommend a combination in the case of an MTD not existing for this dose level of agent 2. Our design considers only toxicity responses, which can be improved upon by taking into account efficacy when determining the combination to be used in subsequent phases.

When combining several agents, designs developed for single-agent dose-finding trials cannot be applied to combination studies. For instance, the standard algorithm-based method for phase I dose-finding clinical trials in oncology is the so-called '3 + 3' design, which is referred to as 'memory less' because allocation to the next dose level for an incoming group of three patients depends only upon what has happened to the total of three to six patients previously treated at the current dose level [14–18]. This method was not designed for cases in which the full toxicity ordering is unknown. Indeed, the full ordering between

toxicity probabilities of each combination is not fully known. That is why some investigators fix the dose level of one agent when they set up a combination study. They can thereby bring the problem back to one dimension and use single-agent designs. Another approach studied in this paper is to pre-select a subset of combinations whose toxicity ordering is known and then apply a single-agent design to this reduced number of combinations. This approach performs well if the target dose combinations happen to be included in the subset. However, in practice, because of potential interaction between combined drugs, it can never be guaranteed that the target doses are always pre-selected into the subset of combinations to be investigated. If target doses are not in the subset, the trial would completely fail to find the target doses.

The advantage of our method is that it uses a logistic model that can be understood by a broad panel of readers and the performance seems to outperform existing designs. We also added stopping rules to the design in order to accommodate practical issues. The method can be efficient with different cohort sizes. Nevertheless, estimating parameter values and the different probabilities continues to require some advanced computational skills. To overcome this challenge, we have developed executable files that can be used to (1) determine the next combination and the MTD and estimate toxicity probabilities from data in actual trials and (2) perform simulations before setting up a combination dose-finding trial. These files are freely available upon request.

Acknowledgements

This work was partially funded by grants from the Association Nationale de la Recherche et de la Technologie, 'Laboratoires Servier' (CIFRE number 2011/0900), and Denis Diderot University-Paris 7 ('Aide à la mobilité internationale 2012').

REFERENCES

- [1] Zhou Y. Choice of designs and doses for early phase trials. *Fundamental & Clinical Pharmacology* 2004; **18**(3):373–378.
- [2] Chevret S. *Statistical methods for dose-finding experiments*, Statistics in Practice. John Wiley and Sons Ltd.: Chichester, 2006.
- [3] Common terminology criteria for adverse events (CTCAE) and common toxicity criteria (CTC), National Cancer Institute. Rockville, Maryland, USA, 2009.
- [4] Conaway MR, Dunbar S, Peddada SD. Designs for single- or multiple-agent phase I trials. *Biometrics* 2004; **60**:661–669.
- [5] Wages NA, Conaway MR, O'Quigley J. Continual reassessment method for partial ordering. *Biometrics* 2011; **67**(4):1555–1563.
- [6] Thall PF, Millikan RE, Mueller P, Lee SJ. Dose-finding with two agents in phase I oncology trials. *Biometrics* 2003; **59**(3):487–496.
- [7] Wang K, Ivanova A. Two-dimensional dose finding in discrete dose space. *Biometrics* 2005; **61**:217–222.
- [8] Mandrekar SJ, Cui Y, Sargent DJ. An adaptive phase I design for identifying a biologically optimal dose for dual agent drug combinations. *Statistics in Medicine* 2007; **26**:2317–2330.
- [9] Mandrekar SJ, Qin R, Sargent DJ. Model-based phase I designs incorporating toxicity and efficacy for single and dual agent drug combinations: methods and challenges. *Statistics in Medicine* 2010; **29**:1077–1083.
- [10] Yin G, Yuan Y. A latent contingency table approach to dose finding for combinations of two agents. *Biometrics* 2009; **65**:866–875.
- [11] Yin G, Yuan Y. Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 2009; **58**:211–224.
- [12] Bailey S, Neuenschwander B, Laird G, Branson M. A Bayesian case study in oncology phase I combination dose-finding using logistic regression with covariates. *Journal of Biopharmaceutical Statistics* 2009; **19**:469–484.
- [13] Riviere MK, Dubois F, Zohar S. Competing designs for drug combination in phase I dose-finding clinical trials. *Statistics in Medicine* 2014. DOI: 10.1002/sim.6094.
- [14] Faries D. Practical modifications of the continual reassessment method for phase I cancer clinical trials. *Journal of Biopharmaceutical Statistics* 1994; **4**:147–64.
- [15] Goodman SN, Zahurak ML, Piantadosi S. Some practical improvements in the continual reassessment method for phase I studies. *Statistics in Medicine* 1995; **14**(11):1149–1161.
- [16] Korn EL, Midthune D, Chen TT, Rubinstein LV, Christian MC, Simon RM. A comparison of two phase I trial designs. *Statistics in Medicine* 1994; **13**:1799–806.
- [17] O'Quigley J, Zohar S. Experimental designs for phase I and phase I/II dose-finding studies. *British Journal of Cancer* 2006; **94**(5): 609–613.
- [18] Iasonos A, Wilton AS, Riedel ER, Seshan VE, Spriggs DR. A comprehensive comparison of the continual reassessment method to the standard 3+3 dose escalation scheme in phase I dose-finding studies. *Clinical Trials* 2008; **5**(5):465–477.
- [19] Yuan Y, Yin G. Sequential continual reassessment method for two-dimensional dose finding. *Statistics in Medicine* 2008; **27**(27):5664–5678.
- [20] Ivanova A, Wang K. A non-parametric approach to the design and analysis of two-dimensional dose-finding trials. *Statistics in Medicine* 2004; **23**:1861–1870.
- [21] O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics* 1990; **46**:33–48.
- [22] Zohar S, Resche-Rigon M, Chevret S. Using the continual reassessment method to estimate the minimum effective dose in phase II dose-finding studies: a case study. *Clinical Trials* 2013; **10**(3):414–421.
- [23] Gilks WR, Best NG, Tan KKC. Adaptive rejection metropolis sampling within Gibbs sampling. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 1995; **44**:455–472.
- [24] Sweeting MJ, Mander AP. Escalation strategies for combination therapy phase I trials. *Pharmaceutical Statistics* 2012; **11**(3):258–266.
- [25] Thall PF, Simon R. Practical Bayesian guidelines for phase IIB clinical trials. *Biometrics* 1994; **50**(2):337–349.
- [26] Storer BE. Design and analysis of phase I clinical trials. *Biometrics* 1989; **45**:925–37.
- [27] Ivanova A, Montazer-Haghighi A, Mohanty SG, Durham SD. Improved up-and-down designs for phase I trials. *Statistics in Medicine* 2003; **22**(1):69–82.
- [28] Simon R, Freidlin B, Rubinstein L, Arbuck SG, Collins J, Christian MC. Accelerated titration designs for phase I clinical trials in oncology. *Journal of the National Cancer Institute* 1997; **89**(15): 1138–1147.
- [29] Gasparini M, Bailey S, Neuenschwander B. Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 2010; **59**: 543–546.
- [30] Cheung YK, Chappell R. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* 2000; **56**(4): 1177–1182.
- [31] Cheung YK, Thall PF. Monitoring the rates of composite events with censored data in phase II clinical trials. *Biometrics* 2002Mar; **58**(1):89–97.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of this article at the publisher's web site.

Appendix

Gibbs sampling

Gibbs sampling was used to generate a sequence of samples from the joint probability distribution of $\beta = (\beta_0, \beta_1, \beta_2, \beta_3)$. The Gibbs sampler relies on the availability of all complete conditional distributions. We start from an arbitrary point $\beta^{(0)} = (\beta_0^{(0)}, \beta_1^{(0)}, \beta_2^{(0)}, \beta_3^{(0)})$ in the distribution, we sampled in turn in each of the full conditional distribution in updating them as the process goes on.

(1) Initialize $\beta^{(0)} = (\beta_0^{(0)}, \beta_1^{(0)}, \beta_2^{(0)}, \beta_3^{(0)})$

For m from 1 to $M + N$

(2) Sample

- $\beta_0^{(m)}$ from its full conditional distribution $f(\beta_0 | \beta_1^{(m-1)}, \beta_2^{(m-1)}, \beta_3^{(m-1)}, \text{data})$
- $\beta_1^{(m)}$ from its full conditional distribution $f(\beta_1 | \beta_0^{(m-1)}, \beta_2^{(m-1)}, \beta_3^{(m-1)}, \text{data})$
- $\beta_2^{(m)}$ from its full conditional distribution $f(\beta_2 | \beta_0^{(m-1)}, \beta_1^{(m-1)}, \beta_3^{(m-1)}, \text{data})$
- $\beta_3^{(m)}$ from its full conditional distribution $f(\beta_3 | \beta_0^{(m-1)}, \beta_1^{(m-1)}, \beta_2^{(m-1)}, \text{data})$

\Rightarrow return $(\beta^{(M+1)}, \dots, \beta^{(M+N)})$

We chose $M = 2000$ number of “burn-in” iterations to be discarded before convergence.

$\beta^{(M)} = (\beta_0^{(M)}, \beta_1^{(M)}, \beta_2^{(M)}, \beta_3^{(M)})$ converges in distribution to the posterior joint distribution $f(\beta_0, \beta_1, \beta_2, \beta_3 | \text{data})$.

The following $N = 5000$ values retained are then considered as a sample from this distribution.

Adaptive Rejection Metropolis Sampling within Gibbs sampling

Adaptive Rejection Sampling (ARS) can not be used to sample from non log-concave distributions. When this is the case, Gilks, Best and Tan [23] propose to replace the rejection sampling in favor of the Hastings-Metropolis algorithm to update one parameter at a time. But to avoid high probabilities of rejection, they adapted the proposal density to the shape of the full conditional density using ARS. They added to ARS a single Hastings-Metropolis step thus creating ARMS within Gibbs sampling. The detailed method can be found in Gilks [23].

Chapter 5

A Bayesian Dose-finding Design for Clinical Trials Combining a Cytotoxic Agent with a Molecularly Targeted Agent

Background: Cytotoxics and MTAs have different action mechanisms, in a simplistic summary, killing cells for the first and blocking their growth for the second. MTAs have emerged in recent years as another option to cytotoxic treatments. Nevertheless, even if some criterion enables to ascertain the right action of the MTA on its target, it is not immediate that it will necessarily induce the expected activity on the cancer and therefore results in efficacy. Moreover, many cytotoxic agents remain the standard treatment for several types of cancer. Consequently, it will be, in some cases, inefficient and unethical to administer only the new MTA, but it can rather be combined and compared to the standard cytotoxic treatment. In addition, the aim of some MTAs is to act inside the cancer cells in order to inactivate it. These cells do not replicate anymore but they are not killed and stay in the human body and can be still observed on the different medical exams performed on the patient. Therefore, their action can be complementary. Moreover, in general combining several agents enables to skirt some drug resistance. For all these reasons, a new challenge in cancer development is to combine both agents, cytotoxic with MTA.

When combining both agents, a possible synergistic effect on efficacy is expected. As efficacy of the MTA is not monotonic and increasing with the dose contrary to the toxicity but rather increases and plateaus, for the combination of cytotoxic and targeted agent, it is not sufficient to study only the safety as the primary endpoint.

Objective: Therefore, our aim is to propose a phase I/II design to enable to combine a cytotoxic agent with a targeted molecule using the characteristics of each agent.

Method: Our goal is to maximize the efficacy while minimizing the toxicity

under an acceptable threshold. We assume that toxicity is quickly evaluable and used a logistic regression model to evaluate toxicity as a binary outcome. In contrast, we assume that efficacy takes a longer time to evaluate; similarly to survival analysis, we use a proportional hazard model to evaluate the efficacy as a time-to-event outcome and incorporated a plateau point. During the conduct of the trial, using MCMC methods we continuously updated the model estimates and posterior distribution of the toxicity and efficacy probabilities in order to use them to assign the next cohort patients to the estimated optimal combination.

Results: For this design, we have encountered the same issue with plateau estimation and chose to simply estimate the plateau at the dose level with the highest posterior probability. Indeed, due to the highest dimension space, the restricted number of patients, and the ethical constraints and adaptive design that do not enable to explore all dose levels with enough patients, the estimation of the plateau point was difficult and very sensitive. We cannot handle all these tricky points. Thus, we decided to develop a dose-finding design with a plateau for a single-agent in order to propose a more efficient solution in a simpler context. For this paper dealing with combinations, we evaluated our design through a simulation study under various practical scenarios and observed that our design performed well by selecting the optimal combination with a high percentage. Nevertheless, one drawback of our method is that the performance of the design sharply decreases with the number of dose level of the MTA. We did not compare to other methods as, at our knowledge, no design was yet propose in this context. Research still need to be done in this field as combining those agents is the future in cancer development.

Appl. Statist. (2015)

A Bayesian dose finding design for clinical trials combining a cytotoxic agent with a molecularly targeted agent

M.-K. Riviere,

Université Paris 5 and Université Paris 6, and Institut de Recherches Internationales Servier, Suresnes, France

Y. Yuan,

University of Texas M. D. Anderson Cancer Center, Houston, USA

F. Dubois

Institut de Recherches Internationales Servier, Suresnes, France

and S. Zohar

Université Paris 5 and Université Paris 6, France

[Received June 2013. Revised April 2014]

Summary. Novel molecularly targeted agents (MTAs) have emerged as valuable alternatives or complements to traditional cytotoxic agents in the treatment of cancer. Clinicians are combining cytotoxic agents with MTAs in a single trial to achieve treatment synergism and better outcomes for patients. An important feature of such combinational trials is that, unlike the efficacy of the cytotoxic agent, that of the MTA may initially increase at low dose levels and then approximately plateau at higher dose levels as MTA saturation levels are reached. Therefore, the goal of the trial is to find the optimal dose combination that yields the highest efficacy with the lowest toxicity and meanwhile satisfies a certain safety requirement. We propose a Bayesian phase I–II design to find the optimal dose combination. We model toxicity by using a logistic regression and propose a novel proportional hazard model for efficacy, which accounts for the plateau in the MTA dose–efficacy curve. We evaluate the operating characteristics of the proposed design through simulation studies under various practical scenarios. The results show that the design proposed performs well and selects the optimal dose combination with high probability.

Keywords: Combination; Cytotoxicity; Dose finding; Molecularly targeted agent; Phase I–II

1. Introduction

Traditional cytotoxic agents have played important roles in combating cancer. However, after decades of research, it has become difficult to find new cytotoxic agents that are substantially more effective than the existing therapeutic strategies. Recently, novel molecularly targeted agents (MTAs), such as small molecules or monoclonal antibodies, have emerged as alternatives or complements to cytotoxic agents for treating cancer (Le Tourneau *et al.*, 2010, 2011, 2012;

Address for correspondence: Sarah Zohar, Institut National de la Santé et de la Recherche Médicale, Unité Mixte de Service 1138, Team 22, Centre de Recherche des Cordeliers, Université Paris 5 et Université Paris 6, Paris, France.

E-mail: sarah.zohar@inserm.fr

Postel-Vinay *et al.*, 2009). Unlike cytotoxic agents, MTAs modulate specific aberrant pathways in cancer cells, while sparing normal tissue. To take advantage of both types of treatment agent, clinicians are exploring the possibility of combining traditional cytotoxic agents with novel MTAs to achieve treatment synergism and better responses for patients.

This new trend of combining cytotoxic agents with MTAs for treating cancer brings new challenges for early phase dose finding trial design. These challenges arise from the difference in the dose–efficacy curves between the two types of treatment agent. For cytotoxic agents, more is better (i.e. a higher dose yields a greater response) until a dose limiting toxicity level is reached. However, the dose–efficacy relationship of the MTA may not follow a monotonic pattern: the efficacy of the MTA often increases at low dose levels and then plateaus (or approximately plateaus) at higher dose levels once a saturation level in the body has been reached (Le Tourneau *et al.*, 2010; Hoering *et al.*, 2011). Although it is possible that efficacy decreases at higher dose levels, here we focus on the case in which efficacy first increases and then plateaus because such a dose–efficacy relationship is much more commonly encountered in practice.

Consequently, the conventional dose finding paradigm of searching for the maximum tolerated dose is not suitable for combinational trials of a cytotoxic agent with an MTA, and it is imperative to consider efficacy and toxicity simultaneously, with the goal of finding the molecularly targeted optimal dose combination (ODC). This is because, once the MTA dose–efficacy curve reaches a plateau, further increases in the dose of the targeted agent will not yield any therapeutic benefit but will potentially result in greater toxicity (Postel-Vinay *et al.*, 2009). In this paper, the ODC is defined as the most efficacious dose combination that yields the lowest toxicity. As the lowest toxicity can still be excessive, we also require that the ODC satisfies a certain safety requirement, e.g. that the probability of toxicity must be lower than a certain upper bound.

Numerous designs have been proposed to find the maximum tolerated dose for trials combining multiple cytotoxic agents, without considering the efficacy end point. For example, Thall *et al.* (2003) developed a Bayesian approach to identify an entire toxicity ‘contour’ of drug combinations. Conaway *et al.* (2004) proposed a dose finding method that was based on the simple and partial orders of drug combinations. Yuan and Yin (2008) proposed a sequential dose finding design that allows single-agent dose finding methods to be used in multiple-agent combination trials. Braun and Wang (2010) proposed a hierarchical-model-based approach for dose finding. Yin and Yuan (2009) developed a Bayesian dose finding method based on a copula-type regression model. Wages *et al.* (2011) extended the continual reassessment method to two-dimensional dose finding. Recently, several phase I–II drug combination trial designs have been proposed to account for both toxicity and efficacy. Focusing on a combination of cytotoxic agents, Huang *et al.* (2007) proposed a phase I–II design based on the ‘3+3’ type of dose escalation scheme, and Yuan and Yin (2008) developed a model-based approach to accommodate toxicity and efficacy for combination trials. Mandrekar *et al.* (2007) proposed a dose finding design for trials combining two MTAs based on a continuation ratio model for trinary outcomes. Cai *et al.* (2014) proposed a flexible dose finding design for trials combining two MTAs, which used a change point model to reflect that the dose–toxicity surface of combinations may plateau. Despite this rich body of literature, no design is available for clinical trials combining a cytotoxic agent with an MTA, which requires simultaneously accounting for the different behaviours of the cytotoxic agent and the MTA. In addition, the existing phase I–II drug combination designs assume that the efficacy outcome is immediately ascertainable; however, this assumption may not hold in many practical situations because, unlike the toxicity end point, the efficacy end point often requires a relatively long time to assess.

We propose a Bayesian phase I–II design to find the ODC for trials combining a cytotoxic

agent with an MTA. We model efficacy as a time-to-event outcome rather than a binary outcome, thereby eliminating the requirement that the efficacy outcomes of treated patients must be fully evaluated before a new cohort can be enrolled in the trial. To account for the feature of the MTA whereby the dose–efficacy curve may initially increase and then plateau, we incorporate a plateau parameter in the proportional hazard model for time to efficacy. We model the binary toxicity outcome by using a logistic regression model. During the trial, we continuously updated the model estimates and use them to assign patients to the ODC.

The remainder of the paper is organized as follows. In Section 2, we present a case-study that has motivated the methodology proposed. In Section 3, we propose toxicity and efficacy models, and describe a dose finding algorithm to identify the ODC. In Section 4, we present simulation studies to evaluate the operating characteristics of the design proposed and investigate its sensitivity to model specifications. We conclude with a brief discussion in Section 5.

The programs that were used to analyse the data can be obtained from

<http://wileyonlinelibrary.com/journal/rss-datasets>

2. A solid tumour clinical trial

Pishvaian *et al.* (2012) reported a phase I dose finding clinical trial for the combination of imatinib and paclitaxel in patients with advanced solid tumours that are refractory to standard therapy. Imatinib is a tyrosine kinase inhibitor which is used in the treatment of multiple cancers, most notably chronic myelogenous leukaemia. Imatinib works by inhibiting the activity of the BCR-Abl tyrosine kinase enzyme that is necessary for cancer development, thus preventing the growth of cancer cells and leading to their death by apoptosis. Because the BCR-Abl tyrosine kinase enzyme exists only in cancer cells and not in healthy cells, imatinib works effectively as an MTA killing only cancer cells through its action. The goal of the trial was to evaluate the safety of combining imatinib with the traditional cytotoxic chemotherapeutic agent paclitaxel, and to determine whether that combination improved the efficacy of imatinib. In the trial, four doses (300, 400, 600 and 800 mg) of imatinib and three doses (60, 80 and 100 mg m⁻²) of paclitaxel were investigated. Most of the grade 3 or 4 toxicities that are related to therapy involved neutropenia, flu-like symptoms and pain. The treatment response was evaluated by using the response evaluation criteria in solid tumours.

This phase I trial adopted the conventional 3+3 design, which unfortunately suffers from several limitations. First, as the 3+3 design requires that the doses under investigation must be monotonically increasing, only a subset of all 12 possible combinations of imatinib and paclitaxel were investigated in the trial. As a result, the trial might not have even examined the most desirable dose in the 4 × 3 dose combination space. Specifically, the trial selectively investigated six dose combinations: (paclitaxel, imatinib) = (60, 300), (60, 400), (80, 400), (80, 600), (100, 600), (100, 800). The original protocol involved an intensive dose schedule, with continuous daily oral administration of imatinib and weekly paclitaxel infusions. However, after treating patients at the first two doses, the regimen resulted in an excessive number of adverse events, and thus the protocol was amended to a less intensive schedule, with intermittent dosing of imatinib. The second limitation of this use of the 3+3 design is that it ignores the important fact that the efficacy of imatinib does not monotonically increase with the dose, and that the maximum tolerated dose may not be the optimal dose for treating patients. Druker (2002) pointed out that, for treating chronic myelogenous leukaemia, a dose of 400–600 mg of imatinib reached the plateau of the dose–response curve. As a result, 400 mg or 600 mg is the dose of imatinib that is commonly used in clinical practice. This result was confirmed in a meta-analysis (Gaftor-Gvili

et al., 2011) of phase III randomized trials, in which no treatment difference was found between 400 mg and higher doses of imatinib. This trial example demonstrates the need for a new dose finding design to handle the clinical trials that combine an MTA with a traditional cytotoxic agent. We apply our design to the trial in Section 4.

3. Methods

3.1. Toxicity model

Consider a trial combining J doses of cytotoxic agent A with K doses of MTA B and denote (j, k) as the combination of the j th dose level of agent A with the k th dose level of agent B. We assume that toxicity (i.e. the dose limiting toxicity defined by the trial investigator) is quickly ascertainable and monotonically increases with the doses of both agents A and B; this assumption generally holds for cytotoxic agents and is plausible for most MTAs.

Let y_i denote the binary toxicity outcome of patient i with $y_i = 1$ indicating a toxicity response, and p_{jk} denote the toxicity probability of combination (j, k) for $j = 1, \dots, J$ and $k = 1, \dots, K$. We model toxicity by using a logistic model as follows:

$$\text{logit}(p_{jk}) = \beta_0 + \beta_1 u_j + \beta_2 v_k \quad (1)$$

where β_0, β_1 and β_2 are unknown parameters, and u_j and v_k are ‘effective’ doses that are ascribed to the j th dose level of agent A and the k th dose level of agent B on the basis of the prior estimates of the single-agent toxicity probabilities for these two dose levels. The procedure of determining the values of u_j and v_k will be described in Section 3.4. We require $\beta_1 > 0$ and $\beta_2 > 0$ so that toxicity monotonically increases with the dose levels of both agents A and B.

Assuming that during the trial conduct, among the n_{jk} patients who are administered the combination (j, k) , m_{jk} patients experienced toxicity, then the likelihood of the toxicity data $\mathcal{D}_{\text{tox}} = \{n_{jk}, m_{jk}\}$ is

$$L(\mathcal{D}_{\text{tox}} | \beta_0, \beta_1, \beta_2) \propto \prod_{j=1}^J \prod_{k=1}^K p_{jk}^{m_{jk}} (1 - p_{jk})^{n_{jk} - m_{jk}}.$$

Letting $\pi(\beta_0, \beta_1, \beta_2)$ denote the prior distribution of β_0, β_1 and β_2 , the posterior is then given by

$$f(\beta_0, \beta_1, \beta_2 | \mathcal{D}_{\text{tox}}) \propto \pi(\beta_0, \beta_1, \beta_2) L(\mathcal{D}_{\text{tox}} | \beta_0, \beta_1, \beta_2). \quad (2)$$

In model (1), we do not include an interactive effect of the two agents (e.g. an interaction term $\beta_3 u_j v_k$) because the reliable estimation of such an interaction term requires a large sample size (e.g. a few hundred), which is typically not available in early phase trials. Our numerical study suggests that including the interaction term does not improve but often impairs the performance of the design (the results are not shown). For dose finding, our goal is not to model the entire dose–toxicity surface accurately, but to obtain an adequate local fit to facilitate dose escalation and de-escalation. A model may provide a poor global fit for the entire dose–toxicity surface; however, as long as the model provides a good local fit around the current combination, it will lead to correct decisions of dose escalation and dose selection (O’Quigley and Paoletti, 2003).

3.2. Efficacy model

Unlike toxicity, which often can be observed quickly, the efficacy response may require a relatively long follow-up time to be scored. In this circumstance, the conventional approach of

treating efficacy as a binary outcome causes a serious logistic issue: when a new patient is enrolled and is waiting for dose assignment, some of the patients who have already been treated in the trial might not have finished their evaluation yet, and thus their response outcomes are not available to make the decision of dose assignment for the new patient. To overcome this difficulty, we herein model the response as a time-to-event outcome, in which the data of the incomplete efficacy evaluations are naturally incorporated in the decision making of dose assignment as censored observations.

Let t denote the time to response. In early phase clinical trials, the typical way to evaluate efficacy is to follow each patient for a fixed period of time T , e.g. 3 months, after the initiation of the treatment. Within the assessment window $(0, T]$, if the patient responds favourably to the treatment (i.e. $t \leq T$), it is scored as a response and otherwise as a non-response. The efficacy of the drug is defined as the response rate at T . Patients' outcomes after T will not be used to define the efficacy of the drug and to make the decision of dose escalation and selection. In other words, the time to response t is always administratively censored at T . Although we cannot observe any t beyond the time point T , it does not cause any issue here because for evaluating the efficacy of the drug and finding the optimal dose, by definition, we are concerned only with the response rate at T , i.e. $1 - S(T)$, where $S(\cdot)$ denotes the survival function of t . For the same reason, conceptually, we can regard $t = \infty$ for the patients who do not benefit from the treatment without affecting the dose finding. A special feature of the trial combining an MTA with a cytotoxic agent is that the dose–efficacy curve behaves differently with respect to the two agents: efficacy is expected to increase monotonically with the dose of the cytotoxic agent, but not with the dose of the MTA. Efficacy often initially increases and then plateaus with the dose of the MTA after the MTA reaches a level of saturation. Let $\lambda_{jk}(t)$ denote the hazard function that is associated with combination (j, k) , and $\mathbb{1}(\mathcal{C})$ denote the indicator function, which takes a value of 1 if \mathcal{C} is true. We model the time to efficacy for the combination of an MTA and a cytotoxic agent by using a proportional hazard model, augmented with a plateau parameter τ , as follows:

$$\lambda_{jk}(t) = \lambda_0(t) \exp[\gamma_1 w_j + \gamma_2 \{z_k \mathbb{1}(k < \tau) + z_\tau \mathbb{1}(k \geq \tau)\}],$$

where $\lambda_0(t)$ is the baseline hazard, and w_j and z_k are effective doses ascribed to the j th dose level of agent A and the k th dose level of agent B on the basis of the prior estimates of the single-agent efficacy probabilities for these two doses, which will be described in the next section. We assume that $\gamma_1 > 0$ and $\gamma_2 > 0$, and therefore efficacy monotonically increases with the dose of the cytotoxic agent A (i.e. w_j). The plateau parameter τ is an integer between 1 and K and indicates at which dose level of agent B (i.e. the MTA) efficacy reaches a plateau. When the dose level is lower than τ , the efficacy monotonically increases with the dose of the MTA (i.e. z_k) through the covariate effect $\gamma_2 \{z_k \mathbb{1}(k < \tau) + z_\tau \mathbb{1}(k \geq \tau)\} = \gamma_2 z_k$, and, when the dose level is equal to or higher than τ , the efficacy plateaus (with respect to the dose level of agent B) with a constant dose effect $\gamma_2 z_\tau$.

Owing to the small sample size of early phase trials, we take a parameter approach and assume an exponential distribution for the time to efficacy with a constant baseline hazard, i.e. $\lambda_0(t) = \lambda_0$, resulting in the following survival function for the time to efficacy:

$$S_{jk}(t) = \exp(-\lambda_0 t \exp[\gamma_1 w_j + \gamma_2 \{z_k \mathbb{1}(k < \tau) + z_\tau \mathbb{1}(k \geq \tau)\}]).$$

Then, the response rate at the end of T for patients who are treated at the combination (j, k) , denoted by q_{jk} , is given by $q_{ij} = 1 - S_{jk}(T)$. In our design, q_{jk} will be used as the measure of efficacy for determining the dose transition and selection.

For patient i , let s_i denote the actual follow-up time, t_i denote the time to response and (j_i, k_i) denote the combination that is administered to the patient. Define $x_i = \min(T, s_i, t_i)$ and censoring indicator $\delta_i = \mathbb{1}(x_i = t_i)$. Given the efficacy data $\mathcal{D}_{\text{eff}} = \{x_i, \delta_i\}$ obtained from n patients,

the likelihood is given by

$$L(\mathcal{D}_{\text{eff}}|\lambda_0, \gamma_1, \gamma_2, \tau) \propto \prod_{i=1}^n \lambda_{j_i k_i}^{\delta_i}(x_i) S_{j_i k_i}(x_i),$$

and the posterior is

$$f(\lambda_0, \gamma_1, \gamma_2, \tau|\mathcal{D}_{\text{eff}}) \propto \pi(\lambda_0, \gamma_1, \gamma_2, \tau) L(\mathcal{D}_{\text{eff}}|\lambda_0, \gamma_1, \gamma_2, \tau), \quad (3)$$

where $\pi(\lambda_0, \gamma_1, \gamma_2, \tau)$ is the prior for the unknown parameters.

3.3. Specification of prior and effective doses

We first discuss the specification of priors for the model parameters. For the toxicity model, we adopted a vague normal prior $N(0, 100)$ for the intercept β_0 , and, following Chevret (1993), we assigned the slopes β_1 and β_2 independent exponential distributions with a rate parameter of 1, i.e. $\beta_1, \beta_2 \sim \text{Exp}(1)$. For the efficacy model, we took vague priors $\lambda_0 \sim \text{Exp}(0.01)$ and $\gamma_1, \gamma_2 \sim \text{Exp}(0.1)$, and assigned τ a multinomial distribution with probability parameters $\pi = (\pi_1, \dots, \pi_K)$, where π_k is the prior probability that the dose–efficacy curve plateaus at dose level k of the MTA. When there is rich information on the location of τ , for example we know the saturation dosage of the MTA from pharmacokinetic and pharmacodynamic studies, we can choose a set of π to reflect the likelihood of each dose level being the plateau point. When there is no good prior information regarding the location of τ , we recommend assigning τ an increasing sequence of prior probabilities (i.e. $\pi_1 < \pi_2 < \dots < \pi_K$) rather than a non-informative flat prior $\pi_1 = \pi_2 = \dots = \pi_K$. This recommendation is based on our experience with numerical studies, in which we found that using a non-informative prior often caused the dose selection to remain at low dose levels owing to the sparsity of data, whereas the prior with increasing π_k s encourages the dose finding algorithm to explore higher dose levels of agent B and actively to learn the shape of the dose–efficacy curve, thereby improving the ODC selection accuracy. In our simulation study, we took $\pi = (0.14, 0.20, 0.28, 0.39)$, which led to good operating characteristics across various scenarios. A summary of prior distributions is given in Table 1. After specifying the prior distributions, we sampled posterior distributions (2) and (3) by using the Gibbs sampler.

We next discuss how to specify the effective doses (i.e. u_j s and v_k s in the toxicity model, and w_j s and z_k s in the efficacy model) on the basis of the prior estimates of the single-agent toxicity and efficacy probabilities. In practice, before two agents are to be combined, each of them typically has been studied individually. For example, before the solid tumour clinical trial that combines imatinib with paclitaxel (Pishvaian *et al.*, 2012), many phase I and II trials have been conducted to study the single-agent toxicity and efficacy profiles for imatinib (Ramanathan

Table 1. Prior distributions for model parameters

<i>Parameter</i>	<i>Prior distribution</i>
β_0	$N(0, 100)$
β_1, β_2	$\text{Exp}(1)$
λ_0	$\text{Exp}(0.01)$
γ_1, γ_2	$\text{Exp}(0.1)$
τ	multinomial(π_1, \dots, π_K) with $\pi_1 < \dots < \pi_K$

et al., 2008; Gibbons *et al.*, 2008; Lipton *et al.*, 2010; van Oosterom *et al.*, 2001) and paclitaxel (Kato *et al.*, 2011; Tsimberidou *et al.*, 2011; Takano *et al.*, 2002; Lim *et al.*, 2010; Horiguchi *et al.*, 2009). Therefore, we often have good prior estimates of the single-agent toxicity and efficacy probabilities for each of the agents. The purpose of defining and using the effective doses is to match the prior estimates of (single-agent) toxicity and efficacy probabilities under our model with those elicited from the prior information. By doing so, we incorporate the available single-agent dose–toxicity and dose–efficacy information in our model and thus improve the efficiency of the design. This approach has been previously used for dose finding in single-agent trials (Chevret, 2006; Zohar *et al.*, 2013) and drug combination trials (Liu and Ning, 2013).

Specifically, let \hat{p}_{j0} and \hat{p}_{0k} denote the estimates of the single-agent toxicity probabilities for the j th level of agent A and the k th level of agent B respectively, and let $\hat{q}_{j0} \equiv 1 - \hat{S}_{j0}(T)$ and $\hat{q}_{0k} \equiv 1 - \hat{S}_{0k}(T)$ denote the estimates of the single-agent efficacy probabilities for the j th level of agent A and the k th level of agent B (at the end of follow-up). Under toxicity model (1), by setting the dosage of agent B (or A) as 0, we obtain the single-agent toxicity model $\text{logit}(p_{j0}) = \beta_0 + \beta_1 u_j$ for agent A and $\text{logit}(p_{0k}) = \beta_0 + \beta_2 v_k$ for agent B. Therefore, on the basis of the prior estimates \hat{p}_{j0} and \hat{p}_{0k} , we backsolve the effective doses u_j and v_k as

$$u_j = \{\text{logit}(\hat{p}_{j0}) - \hat{\beta}_0\} / \hat{\beta}_1,$$

$$v_k = \{\text{logit}(\hat{p}_{0k}) - \hat{\beta}_0\} / \hat{\beta}_2,$$

where $\hat{\beta}_0$ and $\hat{\beta}_1$ are prior means of β_0 and β_1 . Similarly, under efficacy survival model (2), the single-agent efficacy model is $S_{j0}(t) = \exp\{-\lambda_0 t \exp(\gamma_1 w_j)\}$ for agent A and $S_{0k}(t) = \exp\{-\lambda_0 t \exp[\gamma_2 \{z_k \mathbb{1}(k < \tau) + z_\tau \mathbb{1}(k \geq \tau)\}]\}$ for agent B. We determine the effective doses

$$w_j = \frac{\log\{-\log(1 - \hat{q}_{j0}) / (\hat{\lambda}_0 T)\}}{\hat{\gamma}_1},$$

$$z_k = \frac{\log\{-\log(1 - \hat{q}_{0k}) / (\hat{\lambda}_0 T)\}}{\hat{\gamma}_2},$$

where $\hat{\lambda}_0$, $\hat{\gamma}_1$ and $\hat{\gamma}_2$ are prior estimates of the corresponding parameters, and $\hat{\tau}$ is the highest dose level.

3.4. Dose finding algorithm

At the beginning of the trial, data are very sparse and the estimates of the toxicity and efficacy models are highly unreliable. To improve the reliability of dose finding, we use a start-up phase to collect some preliminary data before switching to the formal model-based dose finding algorithm.

We adopted a start-up phase that was similar to that proposed by Huang *et al.* (2007), which divides the dose combination matrix into a sequence of zones along the diagonal from low doses to high doses (Fig. 1) and then conducts a 3+3 type of dose escalation across the zones. Specifically, we initiate the start-up phase by treating the first cohort of three patients at the lowest zone, i.e. the lowest combination (1, 1), and then continuously escalate the dose to higher dose zones until we first encounter a zone in which all doses are ‘closed’. Given a dose, if more than one patient were to experience toxicity out of the three or six patients who have been administered that dose, we close the dose and require that all higher doses (i.e. any combination having a higher dose level of A or B or A and B) are automatically closed and not eligible for use in treating future patients in the start-up phase. More precisely, if we close dose combination (j, k) , we also close higher doses $\{(j', k'); j' \geq j \text{ and } k' \geq k\}$. The closed dose combinations

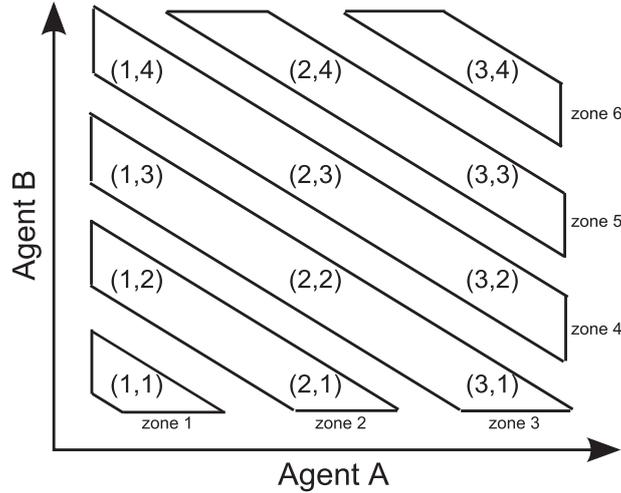


Fig. 1. Illustration of combination zones for the start-up phase

can be reopened later to treat patients in the subsequent model-based dose finding phase if the accumulating data indicate that they are actually safe. The dose escalation across zones is analogous to the traditional 3+3 dose escalation rule: among three patients, if we observe no toxicity, we escalate the dose; if more than two patients experience toxicity, we close the dose; and, if one patient experiences toxicity, we treat three more patients at the current dose. In the last case, if no or one out of the six patients experiences toxicity, we escalate the dose; otherwise we close the dose. When we escalate to a higher dose zone, if there are multiple combinations that are not closed in that zone, we simultaneously assign patients to each of the combinations.

After the start-up phase, we switch to the model-based dose finding phase. Let θ and ξ denote the prespecified toxicity upper bound and efficacy lower bound respectively. Let N denote the total sample size and n denote the number of patients who are treated in the trial. We define that a combination (j, k) is *admissible* if it satisfies the safety requirement

$$P(p_{jk} > \theta) < C_T \tag{4}$$

and also the efficacy requirement

$$P\{S_{jk}(T) > \xi\} \geq C_E \mathbb{1}(n \geq N/2), \tag{5}$$

where C_T and C_E are the respective probability thresholds for toxicity and efficacy. Note that the efficacy requirement (5) takes effect when only half of the patients have been enrolled, as controlled by the indicator function $\mathbb{1}(n \geq N/2)$. We found that introducing the efficacy condition too early caused a high frequency of misclassification of the admissible doses as inadmissible and thus resulted in the early termination of the trial. This situation can arise because, compared with the evaluation of the toxicity condition (4), the reliable evaluation of the efficacy condition (5) requires more data, as the efficacy outcome is not immediately observable and the efficacy model is relatively more complicated.

Let (j, k) denote the current dose, \mathcal{A} denote the set of combinations that have been previously used to treat patients and $\mathcal{B} = \{(j', k'); j' \leq j + 1, k' \leq k + 1, \text{ and } (j', k') \neq (j + 1, k + 1)\}$ denote the set of combinations for which the doses are not two levels higher than the current dose (j, k) . Our model-based dose finding algorithm can be described as follows: after the start-up phase, we assign the next cohort of patients to the optimal combination that is *admissible* and which also has the highest estimate of efficacy, i.e. $1 - \hat{S}(T)$, selected from the set $\mathcal{A} \cup \mathcal{B}$. If several such optimal combinations exist, for example the efficacy has reached a plateau with respect to the

dose level of the MTA, we select the combination with the lowest toxicity probability (e.g. the optimal combination with the lowest MTA dose level) to treat the new cohort. At any time, if all combinations are not admissible, then we terminate the trial; otherwise, we continue this dose assignment process until the maximum sample size is reached. At the end of the trial, we select the ODC as the admissible combination that has the highest estimate of efficacy along with the lowest estimate of toxicity.

4. Numerical studies

4.1. Simulation study

We carried out extensive simulations to evaluate the operating characteristics of the phase I–II design proposed. Taking the setting of the aforementioned solid tumour trial, we assumed three dose levels for cytotoxic agent A (i.e. paclitaxel) and four dose levels for MTA B (i.e. imatinib), resulting in a total of 12 combinations. We took the initial guesses of the single-agent toxicity and efficacy as (0.2, 0.3, 0.4) and (0.3, 0.4, 0.5) respectively for agent A, and (0.12, 0.2, 0.3, 0.4) and (0.3, 0.4, 0.5, 0.59) for agent B. The maximum sample size was 75 and patients were treated sequentially in cohorts of size 3. We assumed that the patient accrual followed a Poisson process with the rate of 1/3.5 patients per week. The toxicity upper bound was $\theta = 0.30$ and the efficacy lower bound was $\xi = 0.20$. We set the toxicity threshold as $C_T = 0.85$ and the efficacy threshold as $C_E = 0.10$, and we took the prior probabilities that the dose–efficacy curve reaches a plateau at the different dose levels of agent B as (0.16, 0.21, 0.27, 0.36). We considered eight different dose–toxicity and dose–efficacy scenarios (see Table 2), representing what we may encounter in practice. We assumed that toxicity was quickly evaluable, whereas the evaluation of efficacy required 7 weeks, i.e. $T = 7$ weeks. Under each scenario, we assumed that, at each combination, the time to efficacy followed an exponential distribution. The parameter of the exponential distribution was chosen such that, at the end of follow-up, the efficacy rate of each dose combination (i.e. $1 - S_{jk}(T)$) matched those displayed in Table 2. As a result, the parameter of the exponential distribution had to vary across doses. Under each scenario, we conducted 1000 simulations.

Table 3 shows the simulation results, including the ODC selection percentage, the average number of patients who were assigned to the ODC and the average number of dose limiting toxicities (DLTs). We also report the ‘effective dose combination’ (EDC) selection percentage, defining the EDCs as the admissible combinations that yield the same (highest) efficacy as the ODC and which also have acceptable toxicity, i.e. the toxicity probability of the EDC is not necessarily the lowest among several equally efficacious combinations. For example, in scenario 2 in Table 2, both (3, 1) and (3, 2) have the same high efficacy rate of 55% and acceptable toxicity rates that are not higher than 30%. The ODC is (3, 1) as it has a lower toxicity probability or lower MTA dose level (i.e. agent B), whereas both (3, 1) and (3, 2) are EDCs as they are both safe and efficacious. Although the ODC is optimal, in practice, the EDCs are also of interest because of their high efficacy even though their dose of agent B may be higher than what is actually needed. Note that, under our definitions, the ODC is one of the EDCs, but not vice versa. Table 4 provides more detailed simulation results for the selection percentages and the number of patients who were treated at each dose combination.

In general, the design proposed performed well across eight scenarios. The ODC and EDC selection percentages were generally greater than 50%, and the design allocated a large number of patients to the ODC and EDCs. Specifically, in scenario 1, the dose–efficacy curve (approximately) plateaus at the lowest dose level of the targeted agent B, and the ODC is the combination (3, 1), which yields the highest efficacy with the lowest dose of the targeted agent. This ODC is also the only EDC in scenario 1. To mimic what may happen in practice, we designed the

Table 2. Eight toxicity and efficacy scenarios for the combination of a cytotoxic agent (agent A) with an MTA (agent B)†

Agent B	Agent A						Agent A					
	1			2			1			2		
	1	2	3	1	2	3	1	2	3	1	2	3
	Toxicity			Efficacy			Toxicity			Efficacy		
	<i>Scenario 1</i>						<i>Scenario 2</i>					
4	0.45	0.50	0.65	0.27	0.42	0.56	0.30	0.45	0.55	0.28	0.43	0.57
3	0.30	0.45	0.55	0.26	0.41	0.56	0.15	0.30	0.45	0.27	0.42	0.56
2	0.15	0.30	0.45	0.25	0.41	0.55	0.12	0.15	<i>0.30</i>	0.26	0.40	<i>0.55</i>
1	0.10	0.15	0.30	0.25	0.40	0.55	0.08	0.10	0.15	0.25	0.40	0.55
	<i>Scenario 3</i>						<i>Scenario 4</i>					
4	0.10	0.15	<i>0.30</i>	0.37	0.52	<i>0.67</i>	0.10	0.15	<i>0.30</i>	0.32	0.46	<i>0.61</i>
3	0.07	0.10	<i>0.15</i>	0.36	0.51	<i>0.66</i>	0.08	0.10	0.15	0.30	0.45	0.60
2	0.05	0.08	0.10	0.35	0.50	0.65	0.04	0.05	0.10	0.20	0.30	0.40
1	0.01	0.05	0.07	0.15	0.25	0.35	0.01	0.03	0.07	0.05	0.10	0.20
	<i>Scenario 5</i>						<i>Scenario 6</i>					
4	<i>0.30</i>	0.55	0.65	<i>0.41</i>	0.56	0.67	0.10	0.15	0.30	0.40	0.55	0.70
3	0.15	0.45	0.55	0.40	0.55	0.65	0.07	0.10	0.15	0.25	0.40	0.55
2	0.10	0.30	0.45	0.15	0.20	0.25	0.05	0.08	0.10	0.15	0.30	0.40
1	0.05	0.15	0.30	0.05	0.10	0.15	0.01	0.05	0.07	0.05	0.10	0.30
	<i>Scenario 7</i>						<i>Scenario 8</i>					
4	0.50	0.60	0.65	0.50	0.60	0.70	0.45	0.55	0.65	0.50	0.60	0.70
3	0.45	0.55	0.60	0.45	0.50	0.60	0.30	0.45	0.60	0.45	0.52	0.63
2	0.30	0.45	0.50	0.30	0.45	0.50	0.15	0.30	0.50	0.30	0.45	0.50
1	0.15	0.30	0.45	0.15	0.30	0.45	0.10	0.15	0.30	0.20	0.30	0.45

†The ODCs are in bold and the effective dose combinations are in italics. The broken lines indicate the dose level of the MTA at which the efficacy plateaus.

scenarios to allow for some variation in efficacy, even when it has reached the plateau. The design proposed selected the ODC 75.6% of the time and allocated on average 35.1 patients to the ODC. As in scenario 1, in scenario 2, the dose–efficacy curve plateaued from the lowest dose level of the targeted agent B with (3, 1) as the ODC, but with two EDCs, i.e. (3, 1) and (3, 2). Combinations (3, 3) and (3, 4) have high efficacy probabilities that are similar to that of the ODC (3, 1), but they are not EDCs because they are not admissible combinations owing to high toxicity. In this case, the design proposed selected the ODC and EDCs 62.0% and 94.0% of the time. In scenario 3, the dose–efficacy curve plateaus after dose level 1 of agent B. The ODC and EDC selection percentages in that scenario were 51.3% and 96.9% respectively. Scenarios 4 and 5 both have efficacy plateaus after dose level 2 of agent B, but with different locations for the ODC and EDCs. In these two cases, the ODC selection percentages were more than 40%.

Table 3. Selection percentage of the ODC and EDC, the average number of patients treated at the ODC and EDC, and the average number of DLTs

	<i>Results (%) for the following scenarios:</i>							
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>
ODC selection percentage	75.6	62.0	51.3	51.3	42.1	66.4	62.8	78.3
Mean number of patients at ODC	35.1	25.1	17.4	18.1	14.4	23.1	32.5	40.9
Mean number of DLTs	22.5	16.7	9.4	10.8	21.9	12.0	25.8	23.7
EDC selection percentage	75.6	94.0	96.9	83.8	68.1	66.4	62.8	78.3
Mean number of patients at EDC	35.1	43.9	42.2	34.4	28.8	23.1	32.5	40.9

Table 4. ODC (**bold**) and EDC (*italic*) selection percentages and the average number of patients treated at each combination

<i>Selection percentage</i>						<i>Average number of patients</i>					
<i>Scenario 1</i>			<i>Scenario 2</i>			<i>Scenario 1</i>			<i>Scenario 2</i>		
0.4	0.1	0.1	0.2	0.2	0.0	2.0	0.2	0.2	3.0	0.9	1.2
0.9	1.0	0.0	0.1	1.0	3.6	4.0	1.3	1.3	3.4	2.7	5.2
0.5	7.4	6.0	0.0	0.8	32.0	3.9	5.8	10.3	3.6	3.5	18.8
0.1	7.9	75.6	0.0	0.0	62.0	4.0	6.9	35.1	3.7	3.8	25.1
<i>Scenario 3</i>			<i>Scenario 4</i>			<i>Scenario 3</i>			<i>Scenario 4</i>		
0.0	0.0	16.7	0.0	0.8	32.5	3.5	2.9	10.3	3.5	3.3	16.3
0.0	0.0	28.9	0.0	0.1	51.3	3.4	3.3	14.5	3.6	3.4	18.1
0.0	0.0	51.3	0.0	0.0	13.9	3.3	3.5	17.4	3.4	3.3	9.4
0.0	0.0	3.1	0.0	0.0	1.4	3.1	3.4	6.5	3.1	3.2	4.5
<i>Scenario 5</i>			<i>Scenario 6</i>			<i>Scenario 5</i>			<i>Scenario 6</i>		
26.0	0.8	0.0	0.0	1.2	66.4	14.4	3.3	0.9	3.4	3.2	23.1
42.1	16.5	1.3	0.0	0.0	19.8	14.4	10.0	2.6	3.5	3.2	13.1
0.1	2.9	3.9	0.0	0.0	9.1	4.4	5.5	5.1	3.4	3.4	7.1
0.0	0.1	5.1	0.0	0.0	3.5	3.4	4.2	6.6	3.1	3.4	5.1
<i>Scenario 7</i>			<i>Scenario 8</i>			<i>Scenario 7</i>			<i>Scenario 8</i>		
1.8	0.0	0.0	4.8	0.3	0.0	2.6	0.2	0.1	4.8	0.7	0.3
8.0	0.1	0.0	16.3	3.1	0.0	7.8	0.7	0.3	9.4	3.3	1.4
27.7	3.6	1.1	3.4	19.5	7.2	14.9	4.1	2.3	5.3	9.5	8.6
2.7	35.1	13.0	0.0	2.8	42.5	7.1	17.6	14.9	3.8	5.7	22.0

Scenarios 6–8 simulate efficacy monotonically increasing with the dose of agent B (for example, agent B does not reach a level of saturation within the range of the investigational doses), with different numbers for the ODCs, which is similar to what may happen in conventional combination trials with two cytotoxic agents. The simulations demonstrate that our proposed design performed well and achieved ODC and EDC selection percentages that were all higher

Table 5. Results of the sensitivity analysis

	<i>Results for the following scenarios:</i>							
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>
<i>Different prior estimates for single agents</i>								
ODC selection percentage	74.7	63.4	50.6	56.9	44.2	65.4	62.2	79.7
Mean number of patients at ODC	74.7	93.6	96.4	85.0	70.0	65.4	62.2	79.7
Mean number of DLTs	0.4	0.0	0.0	0.0	1.8	0.0	7.4	0.3
EDC selection percentage	34.8	25.6	17.1	19.1	15.0	23.0	33.9	41.8
Mean number of patients at EDC	34.8	43.9	41.9	34.4	29.8	23.0	33.9	41.8
<i>Time to efficacy following a Weibull distribution</i>								
ODC selection percentage	78.3	72.0	64.0	47.1	38.1	49.4	64.1	78.5
Mean number of patients at ODC	78.3	96.2	92.6	65.0	56.5	49.4	64.1	78.5
Mean number of DLTs	0.4	0.2	0.0	0.0	0.0	0.0	2.3	0.2
EDC selection percentage	35.7	27.7	20.2	17.3	12.8	16.9	32.8	41.0
Mean number of patients at EDC	35.7	45.2	40.4	29.0	23.5	16.9	32.8	41.0
<i>Double the prior variances of $(\beta_0, \beta_1, \beta_2, \lambda_0, \gamma_1, \gamma_2)$</i>								
ODC selection percentage	70.5	46.3	45.9	52.5	45.9	70.4	61.0	74.4
Mean number of patients at ODC	70.5	88.6	97.9	88.0	72.0	70.4	61.0	74.4
Mean number of DLTs	0.1	0.0	0.0	0.0	0.9	0.0	6.4	0.2
EDC selection percentage	31.5	19.3	16.0	19.1	15.6	24.7	32.3	37.7
Mean number of patients at EDC	31.5	39.8	42.7	35.8	30.1	24.7	32.3	37.7
<i>Different prior for plateau parameter τ</i>								
ODC selection percentage	71.9	54.0	43.3	49.5	42.1	72.4	61.1	78.0
Mean number of patients at ODC	71.9	91.5	96.1	85.8	68.6	72.4	61.1	78.0
Mean number of DLTs	0.3	0.0	0.0	0.0	1.4	0.0	8.3	0.1
EDC selection percentage	32.6	21.7	14.7	18.3	14.7	25.9	32.8	39.5
Mean number of patients at EDC	32.6	42.0	42.6	35.5	29.7	25.9	32.8	39.5

than 60%, suggesting that the design proposed can also be applied to the combination of two cytotoxic agents.

4.2. Sensitivity analysis

We performed sensitivity analyses to study the robustness of our design. We varied four factors:

- (a) the prior estimates of the single-agent toxicity and efficacy probabilities;
- (b) the distribution of time to efficacy;
- (c) the variance of the prior distribution for $\beta_0, \beta_1, \beta_2, \lambda_0, \gamma_1$ and γ_2 ;
- (d) the prior distribution for τ .

We assumed that the single-agent toxicity and efficacy probabilities, which were used to determine the effective dose in the toxicity and efficacy model, were (0.06, 0.12, 0.2) and (0.12, 0.2, 0.3) for agent A and (0.06, 0.12, 0.2, 0.3) and (0.4, 0.5, 0.59, 0.67) for agent B. We simulated the time to efficacy from a Weibull distribution with a fixed shape parameter of 3. We chose the scale parameter of the Weibull distribution such that the efficacy probabilities at the end of follow-up, i.e. $1 - S(T)$, matched those given in Table 2. We inflated the prior variances of parameters $\beta_0, \beta_1, \beta_2, \lambda_0, \gamma_1$ and γ_2 by twofold, and used $\pi = (0.11, 0.17, 0.28, 0.44)$ as the prior probabilities of τ . Table 5 shows the results of the sensitivity analyses. We can see that the ODC and EDC selection percentages and the number of patients who were treated at the ODC and

Table 6. Summarized data from the case-study combination clinical trial involving imatinib with paclitaxel†

<i>Cohort</i>	<i>Paclitaxel dose (mg m⁻²)</i>	<i>Imatinib dose (mg)</i>	<i>Number of evaluable patients for DLT</i>	<i>Number of DLTs</i>	<i>Number of evaluable patients for response</i>	<i>Number of responses</i>
1	60	400	3	0	3	1
2	80	400	3	0	3	2
3	80	600	3	0	3	2
4	100	600	3	0	3	1
5	100	800	6	1	4	3

†The number of responses is the addition of stabilities and partial responses.

EDCs are generally similar to those reported in Table 3, which suggests that the design proposed is not sensitive to the aforementioned design factors.

4.3. Application

We retrospectively applied our design to the solid tumour trial. As described previously, the trial selectively studied six dose combinations out of 12 possible combinations. Because the dosing schedule that was used in the original protocol resulted in too many toxicities, the protocol was amended to use a less intensive dose schedule. As a result, five dose combinations were actually used for dose finding under the amended schedule, as shown in Table 6. The window for assessing treatment response was set at $T = 13$ weeks. The trial did not report the time to response; thus, we assumed that it was uniformly distributed within the assessment window. To be consistent with the 3+3 method that was used by the trial, we set $C_T = 0.33$ and $C_E = 0.0$, and forbade skipping untried doses during the dose escalation.

The trial started by treating the first cohort of three patients at the lowest dose combination (60, 400), at which one response and no DLT was observed. On the basis of the data, our method identified dose combination (100, 600) as the ODC, with an estimated response rate of 0.54, and thus recommended dose escalation to (80, 400) for treating the second cohort of patients. Among the three patients who were treated at (80, 400), two responded to the treatment and no DLT was observed. In light of this new information, our method estimated combination (100, 600) as the ODC, with an estimated response rate of 0.56. Accordingly, we escalated the dose to (80, 600) for treating the third cohort of three patients, of which two patients responded to the treatment and no DLT was observed. Our method then escalated the dose and assigned the fourth cohort to dose combination (100, 600), at which we observed two responses and no DLT. At that moment, the estimated ODC was dose combination (100, 600), with the estimated response rate of 0.45. On the basis of this result, our method would retain the current dose and assign the remaining six patients to (100, 600), whereas the 3+3 design dictated a dose escalation to (100, 800). At the end of the trial, our design selected (100, 600) as the ODC, whereas the 3+3 design picked (100, 800). According to the literature, a dose of 600 mg of imatinib reaches the plateau of the dose–response curve and actually is the dosage that has been widely administered to cancer patients in practice. It seems that our design successfully identified that, whereas the 3+3 design might have resulted in overdosing of patients by selecting a dose of 800 mg.

5. Conclusions

We have proposed a Bayesian phase I–II design for trials that combine a cytotoxic agent with an MTA. We assumed that toxicity is quickly evaluable and used a logistic regression model to

evaluate toxicity as a binary outcome. In contrast, we assumed that efficacy takes a relatively long time to evaluate, and correspondingly we used a proportional hazard model to evaluate efficacy as a time-to-event outcome. To account for the characteristic dose–efficacy curve for MTAs, which initially increases and then plateaus, we incorporated a plateau point in the time-to-efficacy model. During the conduct of the trial, we continuously updated the model estimates and used them to assign patients to the ODC. We evaluated our design through a simulation study under various practical scenarios. Our design performed well by selecting the ODC a high percentage of the time.

The design proposed assumes that the treatment response can be observed any time during the follow-up period. For some clinical trials, the response, however, can only be ascertained at the end of follow-up, e.g. when the response is defined as a certain percentage of tumour shrinkage at time T . In these cases, rather than modelling the time to response, we can model the time to disease progression (i.e. the time to non-response), which typically is observable in realtime on the basis of patients' symptoms. The model and design proposed can still be used. We just need to treat t as the time to non-response, and accordingly to estimate the response rate q_{jk} at T by $S_{jk}(T)$, rather than $1 - S_{jk}(T)$.

There are several possible extensions of the proposed design to improve its performance or flexibility further in order to accommodate different clinical applications. For example, rather than modelling toxicity as a binary outcome, we can use measurements of the various grades of toxicity and model it as an ordinary outcome. This approach uses more refined information and can potentially improve the efficiency of the trial design. In addition, when late onset toxicity is of concern, we can model toxicity as a time-to-event outcome, as we have done for efficacy. Lastly, our dose assignment and selection criteria focus on efficacy while controlling for toxicity. In some applications, it may be more appropriate to consider the trade-off between toxicity and efficacy. To accommodate these cases, we can define a utility function for the trade-off between toxicity and efficacy and then use the utility as the criterion for dose assignment and selection.

Acknowledgements

This work was partially funded by grants from the Association Nationale de la Recherche et de la Technologie, 'Laboratoires Servier' Conventions industrielles de formatique par 12 recherche number 2011/0900), and Denis Diderot University–Paris 7 ('Aide à la mobilité internationale 2012'). Yuan's research was partially supported by awards R01 CA154591 and P50 CA098258 from the National Cancer Institute.

References

- Braun, T. M. and Wang, S. (2010) A hierarchical Bayesian design for phase I trials of novel combinations of cancer therapeutic agents. *Biometrics*, **66**, 805–812.
- Cai, C., Yuan, Y. and Ji, Y. (2014) A Bayesian dose finding design for oncology clinical trials of combinational biological agents. *Appl. Statist.*, **63**, 159–173.
- Chevret, S. (1993) The continual reassessment method in cancer phase I clinical trials: a simulation study. *Statist. Med.*, **12**, 1093–1108.
- Chevret, S. (2006) *Statistical Methods for Dose-finding Experiments*. Chichester: Wiley.
- Conaway, M. R., Dunbar, S. and Peddada, S. D. (2004) Designs for single- or multiple-agent phase I trials. *Biometrics*, **60**, 661–669.
- Druker, B. J. (2002) Perspectives on the development of a molecularly targeted agent. *Cancer Cell*, **1**, 31–36.
- Gafer-Gvili, A., Leader, A., Gurion, R., Vidal, L., Ram, R., Shacham-Abulafia, A., Ben-Bassat, I., Lishner, M., Shpilberg, O. and Raanani, P. (2011) High-dose imatinib for newly diagnosed chronic phase chronic myeloid leukemia patients—systematic review and meta-analysis. *Am. J. Hematol.*, **86**, 657–662.
- Gibbons, J., Egorin, M. J., Ramanathan, R. K., Fu, P., Mulkerin, D. L., Shibata, S., Takimoto, C. H., Mani, S., LoRusso, P. A., Grem, J. L., Pavlick, A., Lenz, H. J., Flick, S. M., Reynolds, S., Lagattuta, T. F., Parise,

- R. A., Wang, Y., Murgo, A. J., Ivy, S. P. and Remick, S. C. (2008) Phase I and pharmacokinetic study of imatinib mesylate in patients with advanced malignancies and varying degrees of renal dysfunction: a study by the National Cancer Institute Organ Dysfunction Working Group. *J. Clin. Oncol.*, **26**, 570–576.
- Hoering, A., LeBlanc, M. and Crowley, J. (2011) Seamless phase I-II trial design for assessing toxicity and efficacy for targeted agents. *Clin. Cancer Res.*, **17**, 640–646.
- Horiguchi, J., Rai, Y., Tamura, K., Taki, T., Hisamatsu, K., Ito, Y., Seriu, T. and Tajima, T. (2009) Phase II study of weekly paclitaxel for advanced or metastatic breast cancer in Japan. *Anticancer Res.*, **29**, 625–630.
- Huang, X., Biswas, S., Oki, Y., Issa, J. P. and Berry, D. A. (2007) A parallel phase I/II clinical trial design for combination therapies. *Biometrics*, **63**, 429–436.
- Kato, K., Tahara, M., Hironaka, S., Muro, K., Takiuchi, H., Hamamoto, Y., Imamoto, H., Amano, N. and Seriu, T. (2011) A phase II study of paclitaxel by weekly 1-h infusion for advanced or recurrent esophageal cancer in patients who had previously received platinum-based chemotherapy. *Cancer Chemther. Pharmacol.*, **67**, 1265–1272.
- Le Tourneau, C., Dieras, V., Tresca, P., Cacheux, W. and Paoletti, X. (2010) Current challenges for the early clinical development of anticancer drugs in the era of molecularly targeted agents. *Target Oncol.*, **5**, 65–72.
- Le Tourneau, C., Gan, H. K., Razak, A. R. and Paoletti, X. (2012) Efficiency of new dose escalation designs in dose-finding phase I trials of molecularly targeted agents. *PLOS ONE*, **7**, no. 12, article e51039.
- Le Tourneau, C., Razak, A. R., Gan, H. K., Pop, S., Dieras, V., Tresca, P. and Paoletti, X. (2011) Heterogeneity in the definition of dose-limiting toxicity in phase I cancer clinical trials of molecularly targeted agents: a review of the literature. *Eur. J. Cancer*, **47**, 1468–1475.
- Lim, W. T., Tan, E. H., Toh, C. K., Hee, S. W., Leong, S. S., Ang, P. C., Wong, N. S. and Chowbay, B. (2010) Phase I pharmacokinetic study of a weekly liposomal paclitaxel formulation (Genexol-PM) in patients with solid tumors. *Ann. Oncol.*, **21**, 382–388.
- Lipton, A., Campbell-Baird, C., Harvey, H., Kim, C., Demers, L. and Costa, L. (2010) Phase I trial of zoledronic acid + imatinib mesylate (Gleevec) in patients with bone metastases. *Am. J. Clin. Oncol.*, **33**, 75–78.
- Liu, S. and Ning, J. (2013) A bayesian dose-finding design for drug combination trials with delayed toxicities. *Baysn Anal.*, **8**, 703–722.
- Mandrekar, S. J., Cui, Y. and Sargent, D. J. (2007) An adaptive phase I design for identifying a biologically optimal dose for dual agent drug combinations. *Statist. Med.*, **26**, 2317–2330.
- van Oosterom, A. T., Judson, I., Verweij, J., Stroobants, S., Donato di Paola, E., Dimitrijevic, S., Martens, M., Webb, A., Scot, R., Van Glabbeke, M., Silberman, S. and Nielsen, O. S. (2001) Safety and efficacy of imatinib (STI571) in metastatic gastrointestinal stromal tumours: a phase I study. *Lancet*, **358**, 1421–1423.
- O’Quigley, J. and Paoletti, X. (2003) Continual reassessment method for ordered groups. *Biometrics*, **59**, 430–440.
- Pishvaian, M. J., Slack, R., Koh, E. Y., Beumer, J. H., Hartley, M. L., Cotarla, I., Deeken, J., He, A. R., Hwang, J., Malik, S., Firozvi, K., Liu, M., Elston, B., Strychor, S., Egorin, M. J. and Marshall, J. L. (2012) A Phase I clinical trial of the combination of imatinib and paclitaxel in patients with advanced or metastatic solid tumors refractory to standard therapy. *Cancer Chemther. Pharmacol.*, **70**, 843–853.
- Postel-Vinay, S., Arkenau, H. T., Olmos, D., Ang, J., Barriuso, J., Ashley, S., Banerji, U., De-Bono, J., Judson, I. and Kaye, S. (2009) Clinical benefit in Phase-I trials of novel molecularly targeted agents: does dose matter?. *Br. J. Cancer*, **100**, 1373–1378.
- Ramanathan, R. K., Egorin, M. J., Takimoto, C. H., Remick, S. C., Doroshow, J. H., LoRusso, P. A., Mulkerin, D. L., Grem, J. L., Hamilton, A., Murgo, A. J., Potter, D. M., Belani, C. P., Hayes, M. J., Peng, B. and Ivy, S. P. (2008) Phase I and pharmacokinetic study of imatinib mesylate in patients with advanced malignancies and varying degrees of liver dysfunction: a study by the National Cancer Institute Organ Dysfunction Working Group. *J. Clin. Oncol.*, **26**, 563–569.
- Takano, M., Kikuchi, Y., Kita, T., Suzuki, M., Ohwada, M., Yamamoto, T., Yamamoto, K., Inoue, H. and Shimizu, K. (2002) Phase I and pharmacological study of single paclitaxel administered weekly for heavily pre-treated patients with epithelial ovarian cancer. *Anticancer Res.*, **22**, 1833–1838.
- Thall, P. F., Millikan, R. E., Mueller, P. and Lee, S. J. (2003) Dose-finding with two agents in Phase I oncology trials. *Biometrics*, **59**, 487–496.
- Tsimberidou, A. M., Letourneau, K., Fu, S., Hong, D., Naing, A., Wheler, J., Uehara, C., McRae, S. E., Wen, S. and Kurzrock, R. (2011) Phase I clinical trial of hepatic arterial infusion of paclitaxel in patients with advanced cancer and dominant liver involvement. *Cancer Chemther. Pharmacol.*, **68**, 247–253.
- Wages, N. A., Conaway, M. R. and O’Quigley, J. (2011) Continual reassessment method for partial ordering. *Biometrics*, **67**, 1555–1563.
- Yin, G. and Yuan, Y. (2009) Bayesian dose finding in oncology for drug combinations by copula regression. *Appl. Statist.*, **58**, 211–224.
- Yuan, Y. and Yin, G. (2008) Sequential continual reassessment method for two-dimensional dose finding. *Statist. Med.*, **27**, 5664–5678.
- Yuan, Y. and Yin, G. (2011) Bayesian phase I/II adaptively randomized oncology trials with combined drugs. *Ann. Appl. Statist.*, **5**, 924–942.
- Zohar, S., Resche-Rigon, M. and Chevret, S. (2013) Using the continual reassessment method to estimate the minimum effective dose in phase II dose-finding studies: a case study. *Clin. Trials*, **10**, 414–421.

Chapter 6

Phase I/II Dose-Finding Design for Molecularly Targeted Agent: Plateau Determination using Adaptive Randomization

Background: More MTAs are developed and evaluated in oncology clinical trials today than ever before. These agents are attractive as they target a specific pathway involved in the cancer growth, and often are less toxic than heavy cytotoxic treatment that are not well tolerated by cancer patients. An FDA guidance pointed out that “cancer vaccine trials have used the ‘3 + 3 design’ and the results show that, except in very rare situations, an MTD for a cancer vaccine may not be identified. In these trials, the dose-toxicity curve may be so flat that the highest dose that can be administered is limited by manufacturing or anatomic issues rather than toxicity” [76]. Therefore, regulatory agencies have highlighted the new challenges encountered for this new kind of molecule and have underlined the need to develop and implement new appropriate statistical dose-finding designs.

For this work, after many discussions with oncologists, it emerged that toxicity is assumed either to be very low or comparable to cytotoxic agents, but efficacy is increasing and then plateaus. Indeed, when all the targeted receptors are already binded to the new drug, there is no need to increase the dose level of the drug, as body saturation is reached. This plateau phenomenon has been observed, for example, in the case of imatinib, a tyrosine kinase inhibitor which is used in the treatment of multiple cancers. Imatinib inhibits the activity of the BCR-Abl tyrosine kinase enzyme that is necessary for cancer development, thus preventing the growth of cancer cells and leading to their death by apoptosis. Because the BCR-Abl tyrosine kinase enzyme exists only in cancer cells and not in healthy cells, imatinib works effectively as an MTA killing only cancer cells. Druker [19] pointed out that for treating chronic myelogenous leukemia a dose of 400 to 600 mg of imatinib reached the plateau of the dose-response curve. This result was confirmed in a meta-analysis [25] of phase III randomized trials, in which no treat-

ment difference was found between 400 mg and higher doses of imatinib. The observance of this phenomenon is still limited, as even if some MTAs can be escalated safely due to their low toxicities, the manufacturing of these agents is limited due to the high costs of the molecules which is increasing with the dose.

Only a very small number of dose-finding methods have been proposed to handle the case that the dose-response curve first increases with the dose and then reaches a plateau. Hunsberger et al. [36] proposed a phase II dose-finding design, assuming a linear regression model for the dose-efficacy curve on the last three dose levels explored. If the estimated slope is positive, the dose is escalated, otherwise if the slope is null or negative, the trial is terminated and the dose with the highest response rate is selected as the optimal dose. Ivanova and Xiao [41] have developed a phase II design to determine the minimum effective dose (MED) when it is located at the breaking point of the plateau. Nevertheless, this design assumes that the value of the efficacy probability on the plateau is known.

Method: We have proposed a Bayesian phase I/II dose-finding design for the MTA using a logistic change point model. Our method focuses on selecting the optimal dose, that is the dose associated with the lowest toxicity among those with highest efficacy, rather than the MTD, in order to reduce toxicity at the same level of efficacy. We proposed two allocation methods, one based on adaptive randomization, where the estimated plateau is sampled among all doses with posterior probabilities for the plateau close enough to the maximum; and the other based on difference between efficacy probabilities after performing Bayesian model averaging on the efficacy probabilities (each model corresponding to each possible location of the plateau).

Results: During the development of this method, we have encountered some issues regarding the estimation of the plateau point in our adaptive design. Indeed, it has been recognized in sequential decision making that algorithms which choose each successive action by optimizing a decision criterion can get stuck at a sub-optimal action. This is due to the fact that the algorithm repeatedly select the sub-optimal action and therefore fails obtaining enough data to select a truly optimal action. This problem is sometimes known as the “optimization versus exploration” dilemma [63, 29, 68] and has been recognized in the context of dose-finding clinical trials [4, 71, 59]. To avoid this issue, we used adaptive randomization [71] in our allocation process. Adaptive randomization samples the referred parameter according to estimated probabilities, therefore it enables to use all the accumulated information through the estimated probabilities, but add a necessary randomness.

Both allocation rules give good and similar performance in terms of percentage of correct selection of the optimal dose, but the one based on posterior probabilities of the plateau location seems more robust across scenarios, the PCS are always above 50%. Moreover, it also gives better results in terms of percentage of selection of a correct dose level, that is dose levels with the highest efficacy but not necessarily lowest toxicity under toxicity restrictions.

We have assumed that efficacy was increasing and then plateaus (if reached) according to our discussion with physicians. Nevertheless, for biologic agents it

was rarely but also observed that toxicity initially increases and possibly decreases for higher doses. We felt that our design could be applied to a wide variety of possible dose-response relationships, including non-monotone. Indeed, when the efficacy is increasing and then decreasing with the dose, we can expect, if patients are not allocated to far from the mode of the unimodal relationship, the decrease will be estimated as a plateau at the correct location. Nevertheless, due to the start-up phase, we can also expect that the decrease will be estimated as a plateau before the true location if the efficacy probabilities are similar at those places. We performed simulations on unimodal scenarios and observed that in general the proposed designs seem to perform well.

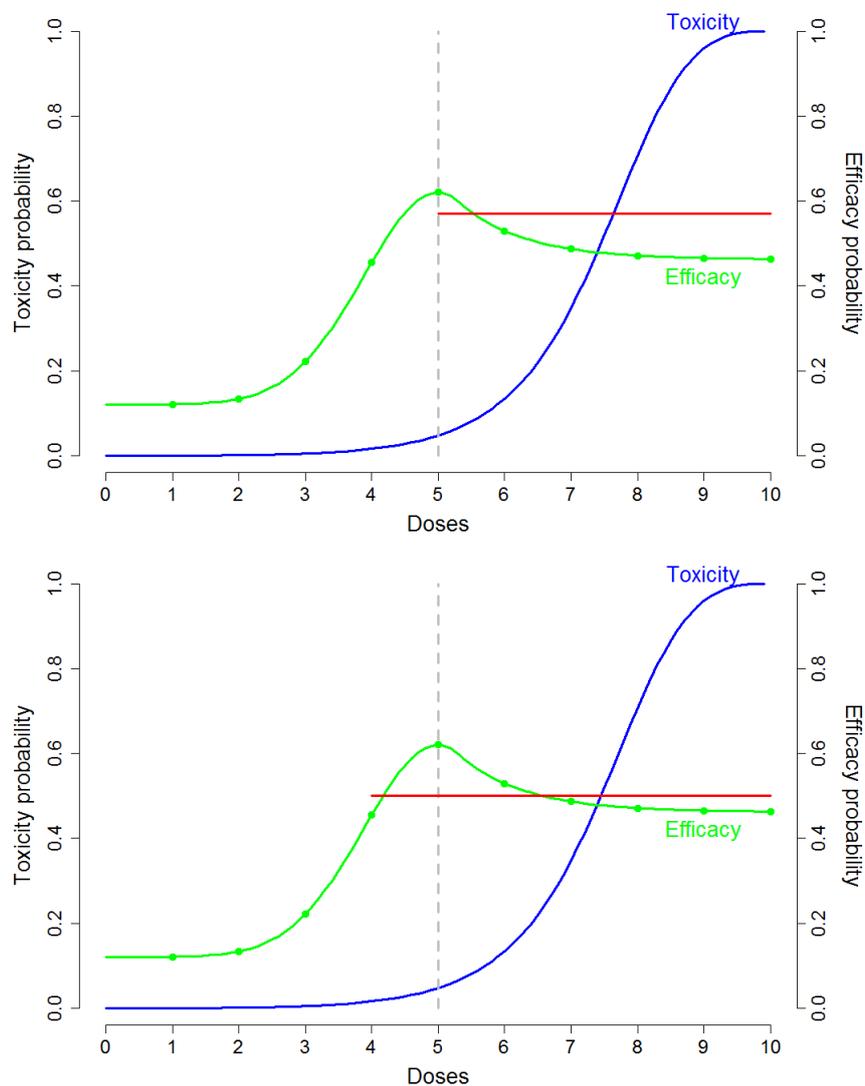


Figure 6.1: Two possible plateau estimations depending on the highest dose explored (not only) when extended our design to unimodal relationships.

Heterogeneity in patients efficacy responses depending on a biomarker can lead to separating the trial into subgroups if there is a strong assumption that the recommended dose level will be different between subgroups. The aim is to provide a treatment adapted to each patient by pursuing personalised healthcare. Indeed, certain biomarkers found on cancer cells can be used to help predict if a certain treatment is likely to work. For instance, in breast cancers, if the cells have too much of a protein called HER2, drugs such as trastuzumab (Herceptin®) can be helpful in treatment, whereas if the cancer cells have normal amounts of HER2, the drugs won't help. Therefore, tumor tissue is checked for biomarker HER2 before treatment is started. In cases where efficacy is not assumed to be the same for each subgroup of patients, it is reasonable to propose a dose allocation study that will recommend different optimal dose level for each subgroup. Nevertheless, in order to better evaluate the safety for which the acceptable toxicity probability is low, in some cases, within the same clinical trial, it may be possible to share the safety data between subgroups while evaluating the efficacy separately. We extend our work to allow sharing some toxicity features between groups in order to get a more accurate estimation regarding the small sample size. The efficacy of the biological agent can be different in groups depending on a biomaker, leading to the recommendation of different best dose level for each subgroup in terms of efficacy in further studies.

For this design, we have tried the use of another plateau selection criterion, the posterior predictive loss (PPL) proposed by Gelfand and Ghosh [27], but results were not satisfying. The PPL is a criterion which minimize the expectation of prediction error. Several conditions on the *loss* function are required. We think that this criterion was not appropriate in the case of binary outcome as our choice of loss function was limited by the required constraints. Nevertheless, with other assumptions, we think that the PPL should be considered and tested.

The following paper is in revision in *Statistics in Medicine*, and a revised version was submitted at the end of August (so recent modifications are still in blue).

Phase I/II Dose-Finding Design for Molecularly Targeted Agent: Plateau Determination using Adaptive Randomization

M-K. Riviere^{a,b†}, Y. Yuan^{c†}, J.-H. Jourdan, F. Dubois^b and S. Zohar^a

Conventionally, phase I dose-finding trials aim to determine the maximum tolerated dose (MTD) of a new drug under the assumption that both toxicity and efficacy monotonically increase with the dose. This paradigm, however, is not suitable for some molecularly targeted agents (MTAs), such as monoclonal antibodies, for which efficacy often increases initially with the dose and then plateaus. For MTAs, the goal is to find the optimal dose, defined as the lowest safe dose that achieves the highest efficacy. We develop a Bayesian phase I/II dose-finding design to find the optimal dose. We employ a logistic model with a plateau parameter to capture the increasing-then-plateau feature of the dose-efficacy relationship. We take the weighted likelihood approach to accommodate the case that efficacy is possibly late-onset. Based on observed data, we continuously update the posterior estimates of toxicity and efficacy probabilities and adaptively assign patients to the optimal dose. The simulation studies show that the proposed design has good operating characteristics. Copyright © 2012 John Wiley & Sons, Ltd.

Keywords: Dose-finding; Phase I; Phase II; Molecularly targeted agents; Oncology

1. Introduction

Traditionally, phase I dose-finding trials aim to determine the maximum tolerated dose (MTD) of a new drug that will be further investigated for efficacy in phase II. This paradigm is built upon the assumption that both toxicity and efficacy monotonically increase with the dose, which is typically true for conventional cytotoxic agents. Recently, molecularly targeted agents (MTAs) have emerged as new therapeutic option in oncology that has changed the practice of cancer patient care. For many MTAs, e.g., monoclonal antibodies, the monotonicity assumption may be violated for efficacy although it typically holds for toxicity. For example, the FDA guidance points out “cancer vaccine trials have used the “3 + 3 design” and the results show that, except in very rare situations, an MTD for a cancer vaccine may not be identified. In these trials, the dose-toxicity curve may be so flat that the highest dose that can be administered is limited by manufacturing or anatomic issues rather than toxicity” [1]. As another example, the efficacy of PTK/ZK (an orally active inhibitor of

^aINSERM, U1138, Equipe 22, Centre de Recherche des Cordeliers, Université Paris 5, Université Paris 6, Paris, France

^bIRIS (Institut de Recherches Internationales Servier), Suresnes, France

^cDepartment of Biostatistics, The University of Texas MD Anderson Cancer Center, Houston, TX 77030, U.S.A.

[†]E-mail: marie-karelle.riviere@crc.jussieu.fr

[†]E-mail: yyuan@mdanderson.org

vascular endothelial growth factor receptor tyrosine kinases) virtually does not change with the dose once it reaches the threshold (or plateau) of 1000mg, which is below the MTD [2, 3]. Further increasing the dose of PTK/ZK to the MTD does not improve its efficacy. As a result, traditional dose-finding methods for cytotoxic agents are not suitable for MTAs because the MTAs do not necessarily need to be administered at their MTDs to achieve maximal efficacy. For MTAs, we are interested in finding the biological optimal dose, which is defined as the lowest safe dose that achieves the highest efficacy, i.e., the dose corresponding to the plateau changing point in the dose-efficacy curve while satisfying certain toxicity requirement (see Figure 1).

A limited number of dose-finding methods have been proposed to handle the case that the dose-response curve first increases with the dose and then reaches a plateau. Hunsberger et al. [4] proposed a phase II dose-finding design, assuming a linear regression model for the dose-response curve. If the estimated slope is null or negative, the trial is terminated and the dose with the highest response rate is selected as the optimal dose. Hirakawa [5] proposed another dose-finding design by jointly modeling a binary toxicity outcome with a continuous efficacy outcome, in which Mahalanobis distance was used to measure the desirability of the dose for dose assignment. Cai, Yuan and Ji [6] proposed a Bayesian phase I/II to handle drug combination trials involving two MTAs when the dose-response curve plateaus at high doses. Ivanova and Xiao [7] have developed a phase II design to determine the minimum effective dose (MED), which does not necessarily correspond to the biological optimal dose.

Our motivating example is the molecularly targeted agent STI571 (imatinib), that acts against the causative molecular event in chronic myeloid leukemia (CML), and was heralded as a major advance in the treatment of CML. Drucker [8] pointed out that “an analysis of responses in blood counts over time suggested that doses of 400 to 600 mg were on the plateau of a dose-response curve”. This was confirmed in a meta-analysis [9] of phase III randomized trials, in which no difference in outcome was found between treatments using the 400 mg dose and those using higher doses. This example supports the assumption of possible efficacy plateau for MTAs, and highlights that the conventional dose-finding paradigm of searching for the MTD is not suitable for MTA trials.

The aim of this paper is to propose a phase I/II dose-finding design for MTAs. We employ a logistic model with a plateau parameter to capture the increasing-then-plateau feature of the dose-efficacy relationship. We take the weighted likelihood approach to accommodate the possibility that efficacy is late-onset. We assume that toxicity monotonically increases with the dose and model it using a logistic model. Based on observed data, we continuously update the posterior estimates of toxicity and efficacy probabilities and adaptively assign patients to the optimal dose. We conduct extensive simulation to examine the operating characteristics of the proposed design.

2. Methods

2.1. Toxicity model

Consider a trial involving K doses of a MTA. We model the toxicity probability of the k^{th} dose level, denoted as ψ_k , using a logistic model as follows:

$$\text{logit}(\psi_k) = \beta_0 + \beta_1 u_k, \quad (1)$$

where β_0, β_1 are unknown parameters, and u_k is the “effective” dose associated with dose level k , which typically differs from the actual clinical dosage. Let $\tilde{\psi}_k$ denote the prior guess of toxicity probability for dose level k , and $\tilde{\beta}_0$ and $\tilde{\beta}_1$ denote the prior estimates of β_0 and β_1 . The “effective” dose u_k is determined by back-solving the dose-toxicity model, i.e.,

$$u_k = \left\{ \log \left(\frac{\tilde{\psi}_k}{1 - \tilde{\psi}_k} \right) - \tilde{\beta}_0 \right\} / \tilde{\beta}_1.$$

The reason that we use the “effective” dose, rather than the actual clinical dosage, in the logistic regression is to regularize the estimate of ψ_k to a practically reasonable range (i.e., *a priori*, the estimate of ψ_k matches our prior guess $\tilde{\psi}_k$), thereby improving the stability of the trial design. This is important because the small sample size of phase I trials often results in very unreliable model fitting, especially at the beginning of the trial when there are only a few observations. Using “effective” dose to fit regression model has been adopted previously in dose-finding studies [10, 11, 12]. In (1), we require $\beta_1 > 0$ so that toxicity monotonically increases with dose levels, as often the case in practice.

Let y_i denote the binary toxicity outcome (1 = toxicity, and 0 = no toxicity) for patients i treated at the dose level $x_i \in \{1, \dots, K\}$, where $i = 1, \dots, N$. After the first I patients are enrolled into the trial, the likelihood of the toxicity data $\mathcal{D}_{\text{tox}} = \{(x_1, y_1), \dots, (x_I, y_I)\}$ is:

$$L(\mathcal{D}_{\text{tox}}|\beta_0, \beta_1) = \prod_{i=1}^I \psi_{x_i}^{y_i} (1 - \psi_{x_i})^{1-y_i}.$$

Letting $f(\beta_0, \beta_1)$ denote the prior distribution of β_0 and β_1 , the posterior is then given by:

$$f(\beta_0, \beta_1|\mathcal{D}_{\text{tox}}) = L(\mathcal{D}_{\text{tox}}|\beta_0, \beta_1)f(\beta_0, \beta_1) \quad (2)$$

We assume prior distributions are independent and take a vague normal prior $N(0, 100)$ for the intercept β_0 , and following [11], we assign the slope β_1 an exponential distribution with a rate parameter of 1, i.e., $\beta_1 \sim \text{Exp}(1)$.

2.2. Efficacy model

A important feature of MTAs that distinguishes them from conventional cytotoxic agents is that the dose-efficacy curves of MTAs do not necessarily increase with the dose. For MTAs, efficacy is often expected to monotonically increase with the dose and then plateau after reaching the level of saturation. Let ϕ_k denote the efficacy probability for dose level k . A general dose-efficacy model can be described as follows,

$$\eta(\phi_k) = g(v_k), \quad (3)$$

where $\eta(\cdot)$ is a link function mapping ϕ_k from $[0, 1]$ to $[-\infty, +\infty]$, and $g(\cdot)$ is a function that plateaus when the effective dose v_k is larger than a certain value. We here adopt the logistic link function with $\eta(\phi_k) = \log\{\phi_k/(1 - \phi_k)\}$, and use a change point model for $g(\cdot)$ as follows:

$$g(v_k) = \gamma_0 + \gamma_1(v_k \mathbb{1}(k < \tau) + v_\tau \mathbb{1}(k \geq \tau)), \quad (4)$$

where $\mathbb{1}(\mathcal{C})$ denote the indicator function, which takes a value of 1 if \mathcal{C} is true, and γ_0 and $\gamma_1 > 0$ are unknown parameters. The plateau parameter τ is an integer between 1 and K that indicates at which dose level the dose-efficacy curve reaches the plateau. When the dose level is lower than τ , efficacy monotonically increases with the dose, and when the dose level is equal to or higher than τ , efficacy plateaus with a constant dose effect $\gamma_1 v_\tau$. We note that other dose-efficacy models can be certainly entertained by choosing different function forms of $\eta(\cdot)$ and $g(\cdot)$. For example, if we take $g(v_k) = E_{\text{max}} \times \frac{v_k}{v_k + ED_{50}}$, with E_{max} denoting the maximum treatment effect and ED_{50} denoting the dose resulting in 50% of E_{max} , we obtain the Emax dose-response model. We here choose to use the change point mode because it explicitly models our dose-finding target, namely the plateau point, through the parameter τ , thereby simplifying the interpretation of the model and decision making of dose assignment. In model (4), the “effective” dose v_k is determined in a similar way as before. That is, we first elicit prior estimates of parameters $\tilde{\phi}_k, \tilde{\gamma}_0, \tilde{\gamma}_1$ and $\tilde{\tau}$ from physicians, and then obtain the value of v_k by back-solving the dose-efficacy model as follows, $v_k = \left\{ \log\left(\frac{\tilde{\phi}_k}{1-\tilde{\phi}_k}\right) - \tilde{\gamma}_0 \right\} / \tilde{\gamma}_1$. In order to make v_k identifiable, we require $\tilde{\phi}_1 < \dots < \tilde{\phi}_K$ (note that this does not mean that the true value of ϕ_k is monotonic). Again, the “effective” dose is not the essential part of the model specification, and the purpose of using the “effective” dose, rather the

actual clinical dosage, is to regularize the prior estimates of efficacy probabilities within a reasonable range and stabilize the estimation of the regression model.

Let z_i denote the binary efficacy outcome (1 = efficacy, and 0 = no efficacy) for patients i treated at the dose level $x_i \in \{1, \dots, K\}$, where $i = 1, \dots, N$. After the first I patients are enrolled into the trial, the likelihood of the efficacy data $\mathcal{D}_{\text{eff}} = \{(x_1, z_1), \dots, (x_I, z_I)\}$ is:

$$L(\mathcal{D}_{\text{eff}}|\gamma_0, \gamma_1, \tau) = \prod_{i=1}^I \phi_{x_i}^{z_i} (1 - \phi_{x_i})^{1-z_i}.$$

Letting $f(\gamma_0, \gamma_1, \tau)$ denote the prior distribution of γ_0, γ_1 , and τ , the posterior is then given by:

$$f(\gamma_0, \gamma_1, \tau|\mathcal{D}_{\text{eff}}) = L(\mathcal{D}_{\text{eff}}|\gamma_0, \gamma_1, \tau)f(\gamma_0, \gamma_1, \tau). \quad (5)$$

We assume prior distributions are independent and took vague normal prior $N(0, 100)$ for the intercept γ_0 and an exponential distributions with a rate parameter of 1 for γ_1 , i.e., $\gamma_1 \sim \text{Exp}(1)$. For the plateau parameter, τ , we assign a discrete prior distribution $\text{pr}(\tau = k) = p_k, k = 1, \dots, K$, with $\sum_{k=1}^K p_k = 1$ and $\forall k, p_k \geq 0$. When there is prior information on the location of τ , e.g., we know the saturation dosage of the MTA from pharmacokinetic and pharmacodynamic studies, we can choose a set of p_k 's to reflect the likelihood of each dose level being the plateau point. When no information is available on the plateau location, the noninformative prior is recommended with $p_1 = \dots = p_K = 1/K$. After specifying the prior distributions, the posterior distribution is sampled using the Gibbs sampler.

Thus far, we have specified the marginal models for toxicity and efficacy. For our purpose of dose finding, these marginal models are adequate because the dose finding algorithm proposed in the following section only relies on the estimates of the marginal probabilities of toxicity and efficacy. Nevertheless, from the theoretical viewpoint, it is appealing to jointly model toxicity and efficacy. To this end, we extend the marginal models by adding a shared random effect b_i to introduce the correlation between toxicity and efficacy as follows:

$$\begin{aligned} \text{logit}(\psi_{k,i}) &= b_i + \beta_0 + \beta_1 u_k \\ \text{logit}(\phi_{k,i}) &= \alpha b_i + \gamma_0 + \gamma_1 (v_k \mathbf{1}(k < \tau) + v_\tau \mathbf{1}(k \geq \tau)), \end{aligned}$$

where b_i follow a normal distribution centered in 0 with variance σ^2 . We assume $1/\sigma^2$ follow a vague prior gamma distribution with parameters (0.1, 0.1). Depending on the value of α , the correlation between toxicity and efficacy can be positive (when $\alpha > 0$), negative (when $\alpha < 0$) or null (when $\alpha = 0$). We choose a prior uniform distribution on $[-10; 10]$ for α . Although this joint model seems preferred from the theoretical viewpoint, numerical studies (described later) show that it does not improve the practical performance of the design. This is because the small sample size of early phase trials contains extremely limited data information for reliably estimating the correlation. Due to this reason, we hereafter mainly focus on the approach based on the marginal models.

2.3. Accommodating delayed efficacy outcome

Unlike toxicity, efficacy often takes a longer follow-up time to assess in practice. Such a ‘‘delayed’’ outcome causes a logistic issue for implementing adaptive designs, that is, when a new patient is enrolled and ready for dose assignment, the patients who have been treated in the trial have not finish their efficacy assessment yet and thus their efficacy outcomes are not available to apply the adaptive rule to assign a dose for the newly accrued patient. Note that we cannot ignore the data information from these partially assessed (or followed) patients because these data are nonignorable, otherwise the resulting estimates are biased [13]. Recently, Liu, Yin and Yuan [14] and Jin et al. [15] proposed a systematical approach to handle delayed outcomes for early phase clinical trials based on Bayesian data augmentation, which enjoys attractive theoretical and practical properties. For simplicity, we here take the approach of Cheung and Chappell [16] by weighting

the observed data likelihood with the follow-up time. Specifically, let T be a fixed time window for evaluating efficacy, and t_i denote the time-to-efficacy of the i^{th} patient. Let $z_{i,I}$ and $C_{i,I}$ denote the response indicator (1 = response, 0 = no response) and the follow-up time for patient i prior to the entry of the $(I + 1)^{\text{th}}$ patient. Before the dose assignment of the $(I + 1)^{\text{th}}$ patient, the weighted likelihood of the efficacy data $\mathcal{D}_{\text{eff}} = \{(x_1, z_{1,I}, w_{1,I}), \dots, (x_I, z_{I,I}, w_{I,I})\}$ obtained from the first I patients is given by

$$L(\mathcal{D}_{\text{eff}} | \gamma_0, \gamma_1, \tau) = \prod_{i=1}^I (w_{i,I} \phi_i)^{z_{i,I}} (1 - w_{i,I} \phi_i)^{1-z_{i,I}},$$

where $w_{i,I}$ takes the form of “adaptive” weights given by

$$w_{i,I} = \begin{cases} 1 & \text{if } t_i \leq C_{i,I} \\ \frac{\#\{j : t_j \leq C_{i,I} \text{ and } C_{j,I} \geq T\} + C_{i,I}/T}{\#\{j : t_j \leq T \text{ and } C_{j,I} \geq T\} + 1} & \text{if } t_i > C_{i,I}. \end{cases}$$

where $\#\{j : t_j \leq T \text{ and } C_{j,I} \geq T\}$ is the number of patients who satisfied the conditions $t_j \leq T$ (i.e., experienced toxicity) and $C_{j,I} \geq T$ (i.e., completed the followup); and $C_{i,I}/T$ is the proportion of the time that patient i was followed compared to the full follow-up time T . Under this weight function, the data (i.e. $z_{i,I}$) from patients whose efficacy outcomes have been observed receive a full weight of $w_{i,I} = 1$. For the patient whose efficacy outcome has not been observed, weight $w_{i,I}$ monotonically increases with the follow-up time $C_{i,I}$. That is, the longer we follow the patient, the more confidence we have about that patient’s current efficacy outcome.

2.4. Dose-finding algorithm

At the beginning of the trial, the posterior estimates of toxicity and efficacy probabilities typically are not reliable due to the limited amount of data [17, 18, 19, 20, 21]. To gather enough information for estimating model parameters, we implement the following start-up phase. Taking a cohort size of 3, we treat the first cohort of patients at the lowest dose level 1, and if no toxicity is observed, we escalate to dose level 2 for treating the second cohort. We continue this dose escalation until encounter the first toxicity. Once a toxicity is observed or the highest dose level is reached, the start-up phase ends and we switch to the model-based dose-finding phase as follows.

Let θ and ξ denote the prespecified toxicity upper bound and efficacy lower bound, respectively. Let $n_{k,I}$ denote the number of patients treated at dose level k , and c denote the cohort size. We define that dose level k is *admissible* if it satisfies the safety requirement

$$P(\psi_k > \theta) < C_T \tag{6}$$

and also the efficacy requirement

$$P(\phi_k > \xi) \geq C_E \mathbb{1}(n_{k,I} > \max(c, 3)), \tag{7}$$

where C_T and C_E are the respective probability thresholds for toxicity and efficacy. Note that the efficacy requirement (7) takes effect only when more than one cohort (or 3 patients) is treated in the trial, as controlled by the indicator function $\mathbb{1}(n_{k,I} > \max(c, 3))$. This is to ensure that we have some data to reliably evaluating the efficacy criterion, given that the efficacy model is relatively complicated.

Let k denote the current dose level and h denote the highest dose level that has been used previously to treat patients, prior to the entry of the $(I + 1)^{\text{th}}$ patient. We use $\mathcal{B} = \{k'; 1 \leq k' \leq \max(\min(k + 1, K), h)\}$ to denote the set of doses that are not one level higher than the current dose k or the highest dose that has been used previously to treat patients, and \mathcal{A} to denote the set of admissible doses in \mathcal{B} . To assign a dose to the incoming $(I + 1)^{\text{th}}$ patient, we fit the proposed model using the data collected from the first I patients enrolled into the trial. Let π_k denote the posterior probability of the k^{th}

dose being the plateau point, i.e., $\Pr(\tau = k | \text{data})$, given by

$$\pi_k = \frac{p_k \iint L(\gamma_0, \gamma_1 | k, \mathcal{D}_{\text{eff}}) f(\gamma_0, \gamma_1) d\gamma_1 d\gamma_0}{\sum_{\tau=1}^K p_\tau \iint L(\gamma_0, \gamma_1 | \tau, \mathcal{D}_{\text{eff}}) f(\gamma_0, \gamma_1) d\gamma_1 d\gamma_0}.$$

Based on the model estimates, we consider two ways of assigning the incoming $(I + 1)^{\text{th}}$ patient (or the new cohort),

- Randomize the $(I + 1)^{\text{th}}$ patient to a dose based on π_k 's. Specifically, let \mathcal{R} denote the set of doses whose posterior probabilities of being the plateau point was close to the largest one with a difference less than a positive threshold s_1 , i.e.,

$$\mathcal{R} = \left\{ j : \left| \max_{1 \leq k \leq K} (\pi_k) - \pi_j \right| \leq s_1; \quad 1 \leq j \leq K \right\}.$$

In other words, \mathcal{R} contains a set of doses that are most likely to be the plateau point. We assign the $(I + 1)^{\text{th}}$ patient to dose $k \in \mathcal{R}$ with a probability $\pi_k / \sum_{j \in \mathcal{R}} \pi_j$. If the randomly selected dose is not in \mathcal{A} (i.e., not admissible), then the patient is assigned to the dose in \mathcal{A} that is closest to that dose. The value of the threshold s_1 should be calibrated by simulation studies. In practice, this can be done as follows: first define a set of representative dose-toxicity scenarios that may be encountered in the trial, and then conduct simulation under different values of s_1 to evaluate the performance of the design. This is a trial-and-error process and may involve repeatedly tuning the values of s_1 based on the simulation results. The goal is to find the values of s_1 that yield good overall performance across different scenarios (e.g., the percentage of correct selection, the number of patients exposed to over-toxic doses or under-toxic doses). Such a calibration based approach has been widely used in clinical trial designs [22, 23, 24, 21]. One version of the threshold we found that generally works well in our simulation study is $s_1 = 0.20 (1 - \frac{I}{N})$. By letting s_1 depend on the current sample size I , the threshold is more liberal at the beginning of the trial when we have high uncertainty on model estimates, and the threshold becomes more stringent toward the end of the trial when we have more data to estimate the model. The above randomization procedure for dose assignment has the advantage of avoiding the dose finding stuck at a suboptimal doses due to high estimation uncertainty.

- Assign the $(I + 1)^{\text{th}}$ patients to the highest admissible dose in \mathcal{A} where we see a big drop on the estimates of efficacy probabilities, i.e., where the dose-efficacy curve is likely to plateau. More precisely, we assign the $(I + 1)^{\text{th}}$ patient to dose

$$k = \underset{\substack{1 \leq j \leq K \\ d_j \in \mathcal{A}}}{\text{argmax}} \left(\hat{\phi}_j - \hat{\phi}_{j-1} \geq s_2 \right),$$

where cutoff s_2 can be interpreted as the minimal efficacy difference of practical importance. The value of s_2 should be calibrated by simulation to ensure good design operating characteristics.

We continue the above dose assignment processes until the maximum sample size is reached. At the end of the trial, we select the optimal dose as the lowest dose level that is admissible and has the highest estimate of efficacy. At any time during the model-based dose-finding phase, if all doses are not admissible, we terminate the trial to protect patients from overly toxic or futile doses.

3. Numerical Studies

3.1. Simulation

We simulated 2000 independent phase I/II trials to evaluate the operating characteristics of the proposed design. We assumed 6 dose levels and considered 10 scenarios (Table 1) with different locations of the true optimal dose. These scenarios cover a wide range of dose-toxicity and -efficacy relationships we may encounter in practice. The prespecified toxicity upper bound and efficacy lower bound was fixed at $\theta = 0.35$ and $\xi = 0.20$, respectively. The maximum sample size was $N = 60$ and the cohort size was $c = 3$ patients. The trial started at the lowest dose d_1 . We took the initial guesses of toxicity and efficacy probabilities as (0.02, 0.06, 0.12, 0.20, 0.30, 0.40) and (0.12, 0.20, 0.30, 0.40, 0.50, 0.59), respectively, to obtain the “effective” doses u_k and v_k used in the toxicity and efficacy models. We set the toxicity threshold as $C_T = 0.90$ and the efficacy threshold as $C_E = 0.40$. We assumed that the patient accrual followed a Poisson process with the rate of 0.28 patients per week, i.e., approximately one patient every 3.5 weeks. We assumed that toxicity was evaluable within 3 weeks, while the evaluation of efficacy required 7 weeks, i.e., $T = 7$ weeks. Under each scenario, we assumed that at each dose, the time to efficacy followed an exponential distribution. The parameter of the exponential distribution was chosen such that at the end of follow-up, the efficacy rate of each dose matched those displayed in Table 1. As a result, the parameter of the exponential distribution had to vary across doses.

We compared the proposed design with the method proposed by Thall and Cook [25] (denoted as the TC design hereafter), which is a phase I/II Bayesian dose-finding design based on trade-offs between the probabilities of efficacy and toxicity. As the TC design assume that the efficacy endpoint is quickly ascertainable, when implementing the TC design, we waited the response of treated patients to completely observed before enrolling a new cohort of patients. The same number of simulations, maximum sample size and cohort size was used. The toxicity and efficacy upper and lower limits were specified as $\bar{\pi}_T = 0.35$ and $\underline{\pi}_E = 0.20$, and p_T and p_E are fixed probability cutoffs both chosen equal to 0.10 as in [25]. For convenience, we refer to the proposed design with the first allocation procedure based on randomization as the MTA-RA, and the proposed design with the second allocation procedure based on posterior mean of efficacy probabilities as the MTA-PM.

As shown in Table 1 and Table 2, the proposed MTA-RA and MTA-PM designs generally perform as well or better than the TC design in terms of both the selection of the optimal dose levels and the number of patients allocated to the optimal doses. In scenario 1, all doses were safe with toxicity probabilities lower than the upper limit, and thus the dose selection was largely guided by efficacy. The percentage of correct selection of the optimal dose (PCS-OD) was 45.0% under the TC design, and greater than 55% under the proposed methods. The MTA-PM design performed best with the highest PCS-OD of 68.7%. In scenarios 2, the plateau was reached at the lowest dose level, the PCS-OD of TC are better than those of the proposed designs in this case. In scenario 3, the third dose is the optimal dose. The PCS-OD of the proposed designs were about 50%, while that of the TC design fall down to 4.5%. The TC design tended to select a higher dose (i.e., the fourth dose) as the optimal dose. Similar results were observed in scenarios 4 to 6. In scenarios 7 and 8, where the efficacy does not exactly plateau but increases by small amount until a difference of 10% from the efficacy plateau, the proposed designs were still able to select the target dose with the highest percentage and substantially outperformed the TC design. In scenarios 9 and 10, none of the doses were admissible and the trial should be terminated. Specifically, in scenario 9, the first three doses was safe but their efficacy was unacceptably low, and the remaining doses were overly toxic; and in scenario 10, all doses were overly toxic. In these two scenarios, the two proposed designs early terminated the trial about than 90% of the time.

In practice, when a dose-finding design missed the optimal dose, but selected a safe dose that has the same efficacy as the optimal dose (although it may have higher dose level than the optimal dose), it is still of interest to physicians. We refers to the set of admissible doses that have the same efficacy as the optimal dose as “correct” doses. For example, in scenario 1, dose level 6 is a “correct” dose, because it has the same efficacy probability of 0.8 as the optimal dose (i.e., dose level 5). Table 2 summarizes the selection percentage of the optimal dose and correct doses under 8 scenarios. We

can see that the selection percentage of the “correct” doses under the proposed designs were mostly over 70%. Between the two proposed design, MTA-RA appeared to have slightly better and more robust performance than the MTA-PM, thus we recommend it for general use in practice.

3.2. Sensitivity analysis

3.2.1. Non-monotonous dose-efficacy relationships

Our designs assume that the dose-efficacy curve increases initially and then plateaus. We conducted a sensitivity analysis to evaluate the performance of the proposed design when the dose-efficacy curve was umbrella shaped (i.e., efficacy first increases and then decreases, see scenarios 1 and 2 in Table 3), monotonically increasing (scenario 3 in Table 3), or monotonically decreasing (scenario 4 in Table 3). The sample size was 36 patients with a cohort size of 3, and we took $\theta = 0.40$ and $\xi = 0.20$. The simulation results (see Table 3) show that the proposed design performed well under these different shapes of dose-efficacy curves. The selection percentages of the target dose were comparable to these reported in Table 1, suggesting that the proposed designs were not sensitive to the violation of the increasing-then-plateau assumption.

3.2.2. Correlation between toxicity and efficacy

We compared the performance of the proposed design using the marginal models, which ignore the correlation between toxicity and efficacy, to the design using the joint model that accounts for the correlation, as described in Section 2.2. Let $\pi_{y,z} = \Pr(y_i = y, z_i = z)$, where $y, z \in \{0, 1\}$. We generated correlated toxicity and efficacy data based on the Gumbel model

$$\pi_{y,z} = \psi^y(1-\psi)^{1-y}\phi^z(1-\phi)^{1-z} + (-1)^{y+z}\psi(1-\psi)\phi(1-\phi)\frac{\exp(\delta) - 1}{\exp(\delta) + 1}$$

where δ is a correlation parameter that controls the correlation between toxicity and efficacy. We set the marginal (toxicity and efficacy) probabilities ψ and ϕ as these displayed in Table 1, and take $\delta = -2, -0.8, 0.8, 2$. Simulation results are provided in Table S1 in supplementary materials. It appears that in general using the joint model did not improve the performance of the design. Instead, it often led to slightly worse performance. For example, for the MTA-RA design, using the joint model yielded 5-20% lower PCS-OD than the design using the marginal models. This is because the small sample size of the trial cannot provide adequate information to reliably estimate the association, and the use of the (more complicated) joint model often results in extra noise and estimation uncertainty.

3.2.3. Link function and prior distributions

We also evaluate the robustness of the proposed designs in terms of the link function and prior distribution. We used a probit link function for the efficacy model and changed the prior distributions of slope parameters β_1 (for toxicity) and γ_1 (for efficacy) from $Exp(1)$ to gamma distribution $Gamma(0.5, 0.5)$. Results are provided in Table S2 in supplementary material. We observe that both proposed methods are generally robust to the choice of the link function and prior distributions. The results are similar to these reported in Table 1.

4. Conclusion

We have developed a phase I/II Bayesian dose-finding design for molecularly targeted agent alone or in combination with a fixed dose of cytotoxic agent. Our design takes into account that the efficacy curve is assumed to plateau. Therefore, our method focus on selecting the optimal dose, that is the dose associated with the lowest toxicity among those with

highest efficacy, rather than the MTD, in order to reduce the toxicity for the same efficacy. We have proposed two different allocation methods based on adaptive randomization with posterior probabilities for the plateau parameter, and on difference between posterior mean of efficacy probabilities according to the plateau parameter. Both allocations give good and similar performance in terms of PCS-OD, but MTA-RA seems more robust across scenarios (always above 50%). Moreover, it also gives better results in terms of percentage of selection of a correct dose level (Table 2), that is dose levels with the highest efficacy but not necessarily lowest toxicity under toxicity restrictions, with high percentages. Then, we have also considered the possibility to extend our design to non-monotone relationships where the mode of the distribution should be selected as the optimal dose. In these cases, our design also gives good performance in general, but MTA-RA performed better in the common case where no plateau is observed across all dose levels. For all these reasons, when a statistician is involved in a clinical trial where a plateau efficacy or a unimodal relationship is expected, we recommend to use the MTA-RA design. [Our program was extended to estimate the optimal doses of several biomarker groups with different efficacy probabilities but with shared toxicity probabilities.](#)

Acknowledgements

Riviere's research was partially funded by grants from the ANRT (Association Nationale de la Recherche et de la Technologie), "Laboratoires Servier" (CIFRE number 2011/0900), Denis Diderot University-Paris 7 ("Aide à la mobilité internationale 2012"). Yuans research was partially supported by Award Number CA154591, 5P50CA098258 and CA016672 from the National Cancer Institute.

References

1. US Department of Health, Human Services F, Drug Administration CfBE, Research. Guidance for Industry, Clinical Considerations for Therapeutic Cancer Vaccines Oct 2011; .
2. Ellis LM. Antiangiogenic therapy: more promise and, yet again, more questions. *J. Clin. Oncol.* Nov 2003; **21**(21):3897–3899.
3. Morgan B, Thomas AL, Dreves J, Hennig J, Buchert M, Jivan A, Horsfield MA, Mross K, Ball HA, Lee L, et al.. Dynamic contrast-enhanced magnetic resonance imaging as a biomarker for the pharmacological response of PTK787/ZK 222584, an inhibitor of the vascular endothelial growth factor receptor tyrosine kinases, in patients with advanced colorectal cancer and liver metastases: results from two phase I studies. *J. Clin. Oncol.* Nov 2003; **21**(21):3955–3964.
4. Hunsberger S, Rubinstein LV, Dancey J, Korn EL. Dose escalation trial designs based on a molecularly targeted endpoint. *Stat Med* Jul 2005; **24**(14):2171–2181.
5. Hirakawa A. An adaptive dose-finding approach for correlated bivariate binary and continuous outcomes in phase I oncology trials. *Stat Med* Mar 2012; **31**(6):516–532.
6. Cai C, Yuan Y, Ji Y. A Bayesian Phase I/II Design for Oncology Clinical Trials of Combining Biological Agents. *Journal of the Royal Statistical Society: Series C* 2013; **Early view article**:DOI: 10.1111/rssc.12039.
7. Ivanova A, Xiao C. Dose finding when the target dose is on a plateau of a dose-response curve: comparison of fully sequential designs. *Pharm Stat* Sep 2013; **12**(5):309–314.
8. Druker BJ. Perspectives on the development of a molecularly targeted agent. *Cancer Cell* Feb 2002; **1**(1):31–36.
9. Gafter-Gvili A, Leader A, Gurion R, Vidal L, Ram R, Shacham-Abulafia A, Ben-Bassat I, Lishner M, Shpilberg O, Raanani P. High-dose imatinib for newly diagnosed chronic phase chronic myeloid leukemia patients—systematic review and meta-analysis. *Am. J. Hematol.* Aug 2011; **86**(8):657–662.
10. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics* 1990; **46**:33–48.
11. Chevret S. The continual reassessment method in cancer phase I clinical trials: a simulation study. *Stat Med* Jun 1993; **12**(12):1093–1108.
12. Liu S, Ning j. Bayesian dose-finding design for drug combination trials with delayed toxicities. *Bayesian Analysis* 2013; **8**:703–722.
13. Yuan Y, Yin G. Robust EM Continual Reassessment Method in Oncology Dose Finding. *Journal of the American Statistical Association* Nov 2011; **106**:818–831.
14. Liu S, Yin G, Yuan Y. Bayesian data augmentation dose finding with continual reassessment method and delayed toxicity. *Ann. Appl. Stat.* 2013; **7**(4):2138–2156.

15. Jin IH, Liu S, Thall PF, Yuan Y. Using data augmentation to facilitate conduct of phase iii clinical trials with delayed outcomes. *Journal of the American Statistical Association* 2014; **109**(506):525–536.
16. Cheung YK, Chappell R. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* Dec 2000; **56**(4):1177–1182.
17. Storer BE. Design and analysis of phase I clinical trials. *Biometrics* 1989; **45**:925–37.
18. Korn EL, Midthune D, Chen TT, Rubinstein LV, Christian MC, Simon RM. A comparison of two phase I trial designs. *Stat Med* 1994; **13**:1799–806.
19. Ivanova A, Montazer-Haghighi A, Mohanty SG, Durham SD. Improved up-and-down designs for phase I trials. *Stat Med* Jan 2003; **22**(1):69–82.
20. Ivanova A, Wang K. A non-parametric approach to the design and analysis of two-dimensional dose-finding trials. *Stat Med* 2004; **23**:1861–1870.
21. Yin G, Yuan Y. A latent contingency table approach to dose finding for combinations of two agents. *Biometrics* 2009; **65**:866–875.
22. Thall PF, Simon R. Practical Bayesian guidelines for phase IIB clinical trials. *Biometrics* Jun 1994; **50**(2):337–349.
23. Thall PF, Millikan RE, Mueller P, Lee SJ. Dose-finding with two agents in Phase I oncology trials. *Biometrics* Sep 2003; **59**(3):487–496.
24. Yin G, Yuan Y. Bayesian dose finding in oncology for drug combinations by copula regression. *JRSS* 2009; **58**:211–224.
25. Thall PF, Cook JD. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* Sep 2004; **60**(3):684–693.

Table 1. Selection percentage and the number of patients (shown in parentheses) allocated to each dose under the TC design and the proposed MTA-RA and MTA-PM designs. In bold are given the optimal doses, and “correct” doses are underlined.

Design	Dose levels						Early termination
	1	2	3	4	5	6	
Scenario 1							
(ψ_k, ϕ_k)	(0.005,0.01)	(0.01,0.10)	(0.02,0.30)	(0.05,0.50)	(0.10,0.80)	(0.15,0.80)	
TC	0.0 (3.2)	0.0 (3.0)	0.0 (3.2)	4.5 (6.9)	45.0 (23.5)	<u>50.5</u> (20.2)	0.0
MTA-RA	0.0 (3.3)	0.1 (3.7)	0.6 (5.0)	4.6 (8.5)	58.5 (21.4)	<u>35.7</u> (17.8)	0.7
MTA-PM	0.2 (3.9)	0.1 (3.7)	7.3 (6.9)	12.3 (9.3)	68.7 (29.8)	<u>10.1</u> (6.0)	1.4
Scenario 2							
(ψ_k, ϕ_k)	(0.01,0.40)	(0.05,0.40)	(0.10,0.40)	(0.25,0.40)	(0.50,0.40)	(0.70,0.40)	
TC	79.5 (48.3)	<u>1.5</u> (1.3)	<u>11.5</u> (4.8)	<u>7.5</u> (4.9)	0.0 (0.5)	0.0 (0.1)	0.0
MTA-RA	63.5 (17.0)	<u>19.3</u> (13.0)	<u>8.3</u> (10.9)	<u>5.4</u> (11.7)	2.7 (6.4)	0.4 (0.8)	0.5
MTA-PM	70.2 (34.4)	<u>18.6</u> (13.3)	<u>8.8</u> (8.0)	<u>1.4</u> (3.0)	0.2 (1.0)	0.0 (0.1)	0.9
Scenario 3							
(ψ_k, ϕ_k)	(0.01,0.25)	(0.02,0.45)	(0.05,0.65)	(0.10,0.65)	(0.20,0.65)	(0.30,0.65)	
TC	37.0 (27.3)	1.0 (2.1)	4.5 (4.2)	<u>38.5</u> (17.6)	<u>18.5</u> (7.3)	<u>1.5</u> (1.4)	0.0
MTA-RA	2.0 (6.8)	14.5 (10.1)	48.3 (14.6)	<u>20.1</u> (12.8)	<u>8.7</u> (10.7)	<u>6.5</u> (5.1)	0.0
MTA-PM	8.9 (8.3)	23.5 (14.2)	49.8 (23.8)	<u>15.6</u> (9.8)	<u>1.8</u> (2.8)	<u>0.0</u> (0.9)	0.5
Scenario 4							
(ψ_k, ϕ_k)	(0.01,0.05)	(0.02,0.25)	(0.05,0.45)	(0.10,0.70)	(0.25,0.70)	(0.50,0.70)	
TC	0.0 (4.5)	0.0 (3.2)	1.5 (4.5)	62.0 (30.0)	<u>35.5</u> (16.2)	1.0 (1.5)	0.0
MTA-RA	0.0 (4.0)	1.0 (5.7)	8.5 (10.0)	53.9 (17.4)	<u>27.5</u> (15.7)	8.9 (7.1)	0.4
MTA-PM	0.3 (4.3)	3.5 (5.6)	22.6 (13.4)	47.0 (22.0)	<u>24.8</u> (12.9)	0.2 (1.0)	1.7
Scenario 5							
(ψ_k, ϕ_k)	(0.01,0.10)	(0.05,0.35)	(0.15,0.60)	(0.20,0.60)	(0.45,0.60)	(0.60,0.60)	
TC	2.0 (7.6)	4.0 (5.3)	30.5 (15.9)	<u>61.5</u> (27.6)	3.0 (3.3)	0.0 (0.3)	0.0
MTA-RA	0.1 (5.5)	8.7 (9.1)	55.4 (17.4)	<u>26.4</u> (15.7)	8.4 (10.2)	1.0 (2.0)	0.2
MTA-PM	1.5 (5.8)	12.8 (10.1)	53.4 (25.6)	<u>27.5</u> (14.7)	2.0 (2.7)	0.0 (0.1)	2.9
Scenario 6							
(ψ_k, ϕ_k)	(0.01,0.05)	(0.05,0.10)	(0.10,0.20)	(0.20,0.35)	(0.30,0.55)	(0.50,0.55)	
TC	0.0 (4.8)	0.5 (3.7)	11.5 (8.3)	37.0 (18.8)	41.5 (17.6)	5.5 (5.2)	4.0
MTA-RA	0.0 (4.7)	1.0 (5.8)	4.7 (8.6)	18.5 (13.0)	55.2 (17.8)	12.7 (7.2)	8.0
MTA-PM	3.2 (8.7)	0.6 (5.7)	8.6 (9.3)	19.9 (11.9)	39.5 (15.5)	2.6 (2.6)	25.7
Scenario 7							
(ψ_k, ϕ_k)	(0.02,0.30)	(0.07,0.50)	(0.13,0.70)	(0.17,0.73)	(0.25,0.76)	(0.30,0.77)	
TC	57.5 (37.2)	2.0 (2.4)	12.0 (7.0)	<u>23.5</u> (10.9)	<u>4.0</u> (2.0)	<u>1.0</u> (0.5)	0.0
MTA-RA	1.4 (6.1)	8.6 (9.0)	38.7 (15.1)	<u>22.9</u> (13.8)	<u>16.6</u> (11.1)	<u>11.8</u> (4.9)	0.0
MTA-PM	10.1 (8.8)	22.9 (15.1)	48.9 (24.6)	<u>16.2</u> (9.3)	<u>1.5</u> (1.7)	<u>0.0</u> (0.4)	0.5
Scenario 8							
(ψ_k, ϕ_k)	(0.03,0.30)	(0.06,0.50)	(0.10,0.52)	(0.20,0.54)	(0.40,0.55)	(0.50,0.55)	
TC	55.5 (35.7)	2.5 (2.6)	<u>14.0</u> (7.2)	<u>25.5</u> (12.4)	2.0 (1.8)	0.0 (0.4)	0.0
MTA-RA	13.5 (10.4)	43.7 (15.0)	<u>20.0</u> (12.8)	<u>12.5</u> (12.0)	8.3 (8.1)	2.0 (1.8)	0.1
MTA-PM	25.5 (16.0)	43.9 (22.4)	<u>24.6</u> (15.1)	<u>5.1</u> (4.8)	0.4 (1.3)	0.0 (0.2)	0.6
Scenario 9							
(ψ_k, ϕ_k)	(0.05,0.01)	(0.10,0.02)	(0.25,0.05)	(0.55,0.35)	(0.70,0.55)	(0.90,0.70)	
TC	0.0 (3.4)	0.0 (3.2)	8.5 (9.5)	9.5 (11.5)	0.5 (1.6)	0.0 (0.3)	81.5
MTA-RA	0.0 (5.8)	0.0 (5.9)	3.0 (7.5)	5.5 (10.9)	0.0 (2.3)	0.0 (0.3)	91.6
MTA-PM	0.1 (6.1)	0.0 (6.0)	2.7 (6.9)	6.6 (10.2)	0.2 (3.2)	0.0 (0.3)	90.4
Scenario 10							
(ψ_k, ϕ_k)	(0.50,0.40)	(0.60,0.55)	(0.69,0.65)	(0.76,0.65)	(0.82,0.65)	(0.89,0.65)	
TC	10.5 (18.7)	2.0 (0.8)	0.0 (0.3)	0.0 (0.1)	0.0 (0.0)	0.0 (0.0)	87.5
MTA-RA	10.9 (16.1)	0.2 (4.0)	0.0 (1.0)	0.0 (0.2)	0.0 (0.0)	0.0 (0.0)	89.0
MTA-PM	10.9 (16.3)	0.1 (4.7)	0.0 (0.7)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	89.0

Table 2. Selection percentages of the optimal dose and “correct” doses without random effect.

Design	Scenarios							
	1	2	3	4	5	6	7	8
	Optimal dose selection percentage							
TC	45.0	79.5	4.5	62.0	30.5	41.5	12.0	2.5
MTA-RA	58.5	63.5	48.3	53.9	55.4	55.2	38.7	43.7
MTA-PM	68.7	70.2	49.8	47.0	53.4	39.5	48.9	43.9
	Correct dose selection percentage							
TC	95.5	100.0	63.0	97.5	92.0	41.5	40.5	42.0
MTA-RA	94.2	96.5	83.6	81.4	81.8	55.2	90.0	76.2
MTA-PM	78.8	99.0	67.2	71.8	80.9	39.5	66.6	73.6

Table 3. Sensitivity analysis of the proposed MTA-RA and MTA-PM designs. In bold are given the optimal doses.

Design	Dose levels								
	1		2		3		4		None
	Scenario 1								
	(0.01, 0.10)		(0.05, 0.35)		(0.15, 0.60)		(0.25, 0.30)		
TC	2.5	(5.0)	5.5	(5.0)	46.5	(12.6)	44.0	(13.3)	1.5
MTA-RA	1.9	(6.3)	26.0	(9.3)	66.4	(13.1)	3.5	(7.0)	2.3
MTA-PM	13.3	(9.3)	25.8	(10.1)	48.8	(13.0)	1.0	(2.5)	11.2
	Scenario 2								
	(0.10, 0.50)		(0.20, 0.70)		(0.30, 0.60)		(0.50, 0.40)		
TC	67.0	(24.8)	26.5	(8.1)	5.5	(2.5)	0.5	(0.6)	0.0
MTA-RA	38.6	(10.8)	51.1	(12.1)	9.3	(9.8)	0.7	(3.2)	0.4
MTA-PM	41.2	(15.9)	56.4	(17.5)	2.3	(2.2)	0.1	(0.4)	0.1
	Scenario 3								
	(0.05, 0.02)		(0.10, 0.28)		(0.16, 0.50)		(0.22, 0.80)		
TC	0.0	(3.3)	2.5	(4.5)	19.0	(8.9)	78.5	(19.3)	0.0
MTA-RA	0.0	(4.4)	3.7	(6.2)	15.0	(9.5)	79.2	(15.5)	2.2
MTA-PM	2.2	(5.9)	12.9	(7.7)	27.8	(11.3)	49.6	(10.1)	7.6
	Scenario 4								
	(0.05, 0.80)		(0.10, 0.50)		(0.16, 0.28)		(0.22, 0.02)		
TC	99.0	(34.9)	1.0	(0.7)	0.0	(0.2)	0.0	(0.1)	0.0
MTA-RA	99.4	(21.1)	0.4	(8.1)	0.0	(5.0)	0.0	(1.2)	0.3
MTA-PM	98.6	(29.8)	1.4	(3.3)	0.0	(1.9)	0.0	(1.0)	0.1

Figure 1. Illustration of dose-toxicity and dose-efficacy relationships for MTAs.

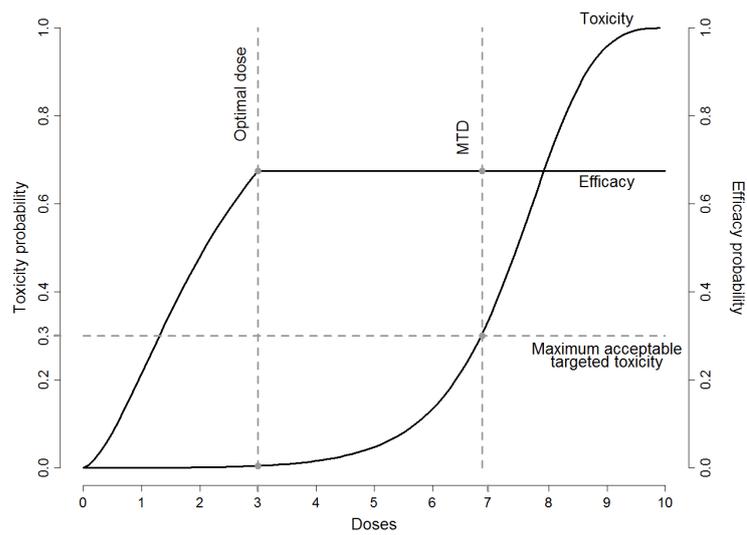


Table S1. Selection percentages of the optimal dose and “correct” doses using random effect model on generated data with correlatio term δ .

$\delta = -2$								
	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8
Optimal dose selection percentage								
MTA-RA	51.7	39.8	36.8	43.5	38.9	42.8	26.6	26.2
MTA-PM	75.8	65.1	43.0	40.6	29.0	40.9	46.2	19.7
Correct dose selection percentage								
MTA-RA	95.1	78.4	93.8	74.6	65.4	42.8	95.8	63.8
MTA-PM	93.7	87.5	77.2	88.4	67.2	40.9	76.0	61.0
$\delta = -0.8$								
	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8
Optimal dose selection percentage								
MTA-RA	51.6	43.8	37.4	43.9	43.5	43.5	28.4	28.7
MTA-PM	77.2	66.2	42.7	41.0	30.7	44.6	47.9	22.5
Correct dose selection percentage								
MTA-RA	94.5	81.8	92.5	74.8	70.9	43.5	96.0	66.6
MTA-PM	94.2	88.1	78.8	89.2	72.2	44.6	78.3	63.0
$\delta = 0.8$								
	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8
Optimal dose selection percentage								
MTA-RA	49.0	43.9	36.4	45.1	42.6	45.2	26.0	29.1
MTA-PM	75.3	66.9	43.9	41.4	28.9	43.3	47.0	21.0
Correct dose selection percentage								
MTA-RA	95.5	83.3	92.3	76.1	70.5	45.2	95.9	68.6
MTA-PM	92.3	88.2	78.4	87.9	71.1	43.3	78.2	60.5
$\delta = 2$								
	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8
Optimal dose selection percentage								
MTA-RA	51.2	43.6	32.9	45.1	39.7	43.5	23.2	28.6
MTA-PM	75.6	68.5	44.2	41.4	25.9	43.4	44.7	20.8
Correct dose selection percentage								
MTA-RA	94.5	82.9	92.7	76.3	68.4	43.5	95.7	70.3
MTA-PM	93.4	87.3	78.6	87.3	67.5	43.4	75.6	61.4

Table S2. Sensitivity analysis with different link functions and prior distributions.

		Probit model for efficacy						
	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8
		Optimal dose selection percentage						
MTA-RA	57.6	61.1	47.5	56.1	55.6	55.1	37.5	44.0
MTA-PM	64.6	76.6	49.6	47.7	56.7	36.2	47.6	39.7
		Correct dose selection percentage						
MTA-RA	93.0	97.4	83.0	82.3	82.2	55.1	88.4	78.6
MTA-PM	77.1	99.1	67.4	70.6	81.0	36.2	66.0	69.1
		Different prior distributions: $\beta_1, \gamma_1 \sim \text{Gamma}(0.5, 0.5)$						
	sc1	sc2	sc3	sc4	sc5	sc6	sc7	sc8
		Optimal dose selection percentage						
MTA-RA	61.1	61.9	45.7	56.0	54.8	50.2	38.2	43.5
MTA-PM	68.5	85.0	45.9	49.3	50.5	40.5	39.7	34.3
		Correct dose selection percentage						
MTA-RA	90.9	96.4	80.8	81.4	78.6	50.2	86.4	74.9
MTA-PM	78.4	99.2	62.8	71.8	76.6	40.5	56.2	57.3

Chapter 7

Conclusion and discussion

A phase I trial is usually the first stage of testing a new drug involving human subjects. Although the treatment has been tested extensively in lab on animals, the side effects in people cannot always be predicted. Therefore, the goal of a phase I trial is to evaluate the safety of the treatment and identify its side effects. Cancer phase I studies are conducted with patients because of the harmfulness of cancer treatments that would not be acceptable on healthy volunteers. Moreover, although these studies are not designed to find out if the new treatment is efficient against cancer, there is a high interest in the new drug's efficacy in those patients directly (considered as a secondary endpoint). These patients often fail standard treatments and are at an advanced stage of the disease. The new drug tested in phase I may be one of their last treatment options. Cancer phase I trials enroll a small number of participants, usually 15 to 50 patients.

The primary aim of phase I clinical trials in oncology is to determine the highest dose level of the new treatment that can be administered with an acceptable toxicity rate, called maximum tolerated dose. The statistical formulation of the problem is to select a dose level from several available doses, with a toxicity probability closest to a given target. Phase I trials are sequential dose-escalation procedures where dose levels of the new drugs are slowly escalated until the observed drug-related toxicity reaches an unacceptable predetermined level.

We performed a systematic review of the literature of all drug combinations phase I trials published the last three years between 1 January 2011 and 31 December 2013, where at least two drugs were planned to undergo dose escalation. Our aim was to determine what were the current practice in combinations trials. Our analysis highlighted that all designs used were for single agent, thus the dose-toxicity relationship was viewed as a one-dimensional dose space while the reality involved several agents inducing a multi-dimensional issue. In particular, most of trials (88%) used a traditional “3+3” or “modified 3+3” design. To bring the problem back into a one-dimensional space, physicians have pre-selected the combinations to be evaluated associated with a known toxicity order. Indeed, most published papers assumed a monotonic and increasing dose-toxicity relationship (62.7% meanwhile 37.3% papers assumed only a partial ordering); this enables the

use of a one-dimensional statistical design instead of multi-dimensional designs. To select the combinations to be considered with a monotonic and increasing dose-toxicity relationship, investigators gradually increase each agent while fixing the others. This process induces a limited number of combinations to be explored and only a subset of combination is evaluated, despite the larger number of possible combinations. Indeed, we observed that the median ratio between the number of combinations considered and the number of possible combination was 0.64, indicating that approximately one third of the combination space was not considered. This means that trial investigators have selected the combinations to be evaluated prior to the trial and some combinations were excluded. To explore the entire combination space is obviously not feasible in practice and physicians may only wish to explore a subset of combinations. Nevertheless, the choice of the combinations to explore should not be limited by partial toxicity ordering and the design should have the possibility to explore any combination it estimates to be the best. Indeed, due to the possible interactions between drugs, pre-selecting an arbitrarily reduced subset of combination induces a risk not to select any combination with a toxicity rate close to the target toxicity. Even if the targeted DLT proportion threshold was 33% in 70% of studies and 16.7% in 7.3% studies, the median toxicity rate associated with the recommended dose at the end of the trial was much lower, in median 5%. Therefore, in general, trials did not manage to achieve the targeted DLT rate. That is maybe a reason why, in 26.4% of the reviewed papers an intermediate combination was added during the course of the trial which induces for some trials a non-monotonic dose-toxicity relationship when one agent is increased while the other is decreased. Therefore, the methods for single agent do not always seem appropriate for combination phase I trials when several agents are varying. These methods are not designed to take into account the multi-dimensional space. Several alternative designs were proposed for combinations either algorithm-based or model-based that give the possibility to explore any appropriate combination in the entire combination space according to the accumulated data. It should be noted that these methods do not allow exploring combinations that are estimated to be too toxic, and they have high operational characteristics. If some combinations are not evaluated during a study this can possibly lead to the failure of some clinical trials to estimate the most suitable combination in terms of toxicity rate. Moreover, in some cases it required to perform amendments in order to explore combinations that were not planned prior to the trial. In the last years, many designs for combination dose-finding studies were proposed that deal with this issue, but they are rarely used in practice, maybe due to a lack of understanding of these designs that require the implication of an expert statistician. Using inappropriate designs, even in early phase clinical trials, can increase attrition rate by proposing a wrong combination choice for phase II and phase III trials.

Building on this finding, a first part of our work consists in the study of several representative designs for combination trials to compare their performance and highlight the advantages or disadvantages of each method. We have chosen two algorithm-based designs and four model-based's. Based on an extensive simulation

study, we have noticed that model-based methods seemed to perform better than algorithm-based methods when targeting a single MTD at the end of the trial. All model-based methods have good operational characteristics with a high percentage of correct selection, and their performances were in general comparable. On this basis, our aim was to propose an innovative adaptive dose-finding design for current practice that (1) would have good operational characteristics in different possible locations of the MTD(s), and (2) would perform, in general, better than existing designs. We decided to model the dose-toxicity relationship of the combination with a logistic model as they are often well-known by physicians and flexible models. We used a 3-parameter logistic regression model with one parameter for each agent and an interaction term. Our dose-allocation process enables the trial to escalate, de-escalate or stay at the same combination depending on the toxicity probability estimation at this dose level and its uncertainty. The MTD recommendation at the end of the trial is based on toxicity intervals. Indeed, the combination selected for further phases was the one with the highest probability for the toxicity rate to be in an interval around the targeted toxicity. According to our simulation study, this method seems to be able to identify the MTD with a high percentage of correct selection in a wide variety of scenarios. We compared our method with other model-based designs for combination drug trials. All the designs seem to be efficient when the MTDs are located on the same diagonal in the combination space. One benefit of our method compared with the other proposed designs is that it is also efficient when the MTDs are not necessarily located on the same diagonal.

For time issues, this design was implemented in C/C++ as it is much faster than R. We are also developing an R package for this method. The package will enable both to perform simulations, and to estimate the next combination in the context of a real clinical trial given the required data. We plan to finish the implementation of this package for the PhD defense.

Over the years, oncologists have prescribed “standard chemotherapy” because they found by trial that these drugs worked well. They reduced the cancer burden in many patients largely by killing rapidly dividing cells. Standard chemotherapy often results in collateral damage to healthy tissue, causing unwanted side effects that impair the circulatory system, the immune system, the digestive system, and others. Because cytotoxic agents usually disrupt molecules and processes that occur in all rapidly dividing cells, many normal cells throughout the body that are undergoing active growth and cell division can also be damaged. Most of the proposed phase I statistical designs that can be found on the literature, as the “3+3”, the CRM or EWOC, have been developed following the assumptions of cytotoxic agents. In recent years, molecularly targeted agent have been developed such as small molecules or antibodies. Unlike standard chemotherapy, targeted therapies are designed to interact with specific molecules that are part of the pathways and processes used by cancer cells to grow, divide, and spread throughout the body. When researchers discover a potentially vulnerable molecule involved in a cancer process or pathway, they design new therapies to disrupt its activity. Many tar-

geted therapies are associated with fewer and less toxic side effects than standard chemotherapy or radiation because they cause little or no collateral damage to normal cells. This can contribute to the quality of life for patients undergoing treatment.

In this context, we have decided to develop a dose-finding design in single-agent for MTAs. This design can be used for MTAs in a combination with a fixed dose of cytotoxic standard therapy which is becoming common in cancer clinical trials. This work was initiated for both a practical need in a real clinical trial inside IRIS with which I am accomplishing my PhD, and discussions with physicians about the difference in assumptions with MTAs. After discussions with physicians and pharmacologists, it emerged that efficacy was assumed to be increasing with the dose and then plateaus. Indeed, when all the targeted receptors were already bound to the new drug, there was no need to increase the dose level of the drug, as body saturation was reached. Further information about the real clinical trial, the molecule or the pathology cannot be detailed for confidentiality reasons. The efficacy was first assumed to be a binary outcome.

We have proposed a Bayesian phase I/II dose-finding design for the MTA using a logistic change point model. Our method focuses on selecting the optimal dose, that is the dose associated with the lowest toxicity among those with highest efficacy, rather than the MTD, in order to reduce the toxicity for the same efficacy. During the development of this method, we have encountered some issues regarding the estimation of the plateau point in our adaptive design. Indeed, it was recognized in sequential decision making that algorithms that choose each successive action by optimizing a decision criterion can get stuck at a sub-optimal action. This is due to the fact that the algorithm repeatedly select the sub-optimal action and therefore fails to obtain enough data and thus to select a truly optimal action. This problem is sometimes known as the “optimization versus exploration” dilemma [63, 29, 68] and has been recognized in the context of dose-finding clinical trials [4, 71, 59]. To avoid this issue, we used adaptive randomization in our allocation process. Adaptive randomization samples the referred parameter according to estimated probabilities, therefore it enables the use of all the accumulated information through the estimated probabilities, but adds a necessary randomness. The two proposed allocation methods give good and similar performance in terms of percentage of correct selection of the optimal dose, but the one based on posterior probabilities of the plateau location seems more robust across scenarios (always above 50%). Moreover, it also gives better results in terms of percentage of selection of a correct dose level, that is dose levels with the highest efficacy but not necessarily lowest toxicity under toxicity restrictions.

Unfortunately, due to population heterogeneity in patients retained for this early phase, efficacy cannot be evaluated accurately and therefore was not considered in the design. The proposed method was finally abandoned and replaced by a CRM. Nevertheless, as the development of MTAs becomes more and more usual practice, another future clinical trial inside the company could be set up using this design. It will be a combination of a targeted therapy with a fixed dose of

standard chemotherapy. Discussions are still on-going, but the same assumptions are made for now.

We also extended this design to (1) unimodal relationships and (2) different biomarker groups leading to different optimal dose in each subgroup with shared toxicity. We observed that, in general, the proposed designs seem to perform well.

For time issues, this design was implemented in C/C++ as it is much faster than R. In order to be easily used in practice, we are currently developing an R package for this method. The package will enable both to perform simulations with a flexible choice of settings, and to estimate the next optimal dose level in the context of a real clinical trial given the required data. We plan to finish the implementation of this package for the PhD defense.

Cytotoxics and MTAs have different action mechanisms, killing cells and blocking their growth by interfering with specific molecules. MTAs have emerged in recent years as another option to cytotoxic treatments. Nevertheless, even if some criterion enables the ascertainment of the right action of the MTA on its target, it is not immediate that it will necessarily induce the expected activity on the cancer and therefore result in efficacy. Moreover, many cytotoxic agents remain the standard treatment of several types of cancer. Consequently, it will be, in some cases, inefficient and unethical to administer only the new MTA, but it can rather be combined and compared to the standard cytotoxic treatment. In addition, the aim of some MTAs is to act inside the cancer cells in order to inactivate it. These cells do not replicate anymore but they are not killed and stay in the human body. They can be still observed on different pathological exams as radiography or MRI (Magnetic Resonance Imaging) performed on the patient, and only PET-Scan (Positron Emission Tomography) can permit to detect their inactivity. Therefore, the action of cytotoxic agents and MTAs can be complementary in order to reduce the spread of the cancer and killing the remaining cancer cells. Furthermore, in general combining several agents enables to skirt some drug resistance. For all these reasons, a new challenge in cancer development is to combine both agents, cytotoxic with MTA. For now, when combining these two agents, investigators choose several dose levels of the new targeted therapy while they are fixing the standard chemotherapy to its approved recommended dose level in single-agent. In this context, single-agent designs for MTAs can be appropriately used. Nevertheless, there is an interest to let both agents vary as the optimal combination in terms of both toxicity and efficacy combined is not necessary when the cytotoxic agent is at its recommended dose in single-agent. When combining several agents, a synergistic effect on efficacy is expected. As efficacy of the MTA is not monotonic and increasing with the dose contrary to the toxicity but rather increases and plateaus, for the combination of cytotoxic and targeted agent, it is not sufficient to study only the safety as the primary endpoint. Therefore, we proposed a phase I/II design to enable to combine a cytotoxic agent with a targeted molecule using the characteristics of each agent. Our goal is to maximize efficacy while minimizing toxicity under an acceptable threshold. We assume that toxicity is quickly ascertainable and we use a logistic regression model to evaluate toxicity

as a binary outcome. In contrast, we assume that efficacy takes a longer time to evaluate, similarly to survival analysis we used a proportional hazard model to evaluate the efficacy as a time-to-event outcome and incorporated a plateau point. During the conduct of the trial, we continuously update the model estimates and posterior distribution of the toxicity and efficacy probabilities in order to use them to assign the next cohort of patients to the estimated optimal combination. For this design, we have encountered the same issue with the plateau estimation and chose to simply estimate the plateau at the dose level with the highest posterior probability. Indeed, due to the higher dimension space, the restricted number of patients, and the use of adaptive designs that do not enable the exploration of all dose levels with sufficient patients, the estimation of the plateau point was difficult and sensitive. It was the result of this work that initiated our decision to develop a dose-finding design with the plateau for single-agent in order to propose a more efficient solution in a simpler context. For this paper dealing with combinations, we evaluated our design through a simulation study under various practical scenarios and observed that our design performed well by selecting the optimal combination with a high percentage. Nevertheless, the performance of the design highly decreases with the number of dose levels of the MTA.

This design was also implemented in C/C++ and as the other methods, the construction of the R package is on-going and should be finished at the end of the PhD.

Perspectives

In perspective of this work, we had two other ideas of project for phase I/II designs of oncology that are presented further.

Moreover, my PhD director, Sarah Zohar, has proposed ideas on another way to develop dose-finding in early phases. It will lead to another PhD she will direct. In phase I oncology clinical trials the toxicity endpoint is usually considered as binary evaluated during a fixed period of time. The common practice of limiting the evaluation window for the occurrence of DLT to only one cycle raises methodological issues regarding; (1) DLTs observed after the evaluation window (2) cumulative toxicities observed over the cycles of treatment. Nevertheless, the dose reduction over several cycles of treatment as well as “late toxicities” observed after the evaluation window, are not taken into account in the final recommendation of the MTD. This means that the dose level recommended at the end of the clinical trial could be far from being the dose level that is given in practice. A practical consequence of this issue is that physicians often have to either reduce the dose level or stop the treatment during the trial (temporarily or permanently). Although doses and protocol standardization guidelines are used by physicians, there are concrete limitations, and the treatment administered to a given patient may be different from the dose recommended in the protocol. Most of the time, non-adherence to protocol is due to drug toxicity. In order to evaluate these deviations from the

protocol and for accurate estimation of the MTD defined in early phase clinical trials, it is important to be able to analyze patient care data. This is especially useful, since (1) phase I patients are different from phase II or III patients, and (2) due to inclusion/exclusion criteria, the subset of patients included in the clinical trials is a subgroup of the treated population in everyday care. A new paradigm of data-driven methodologies reusing healthcare data to provide decision support is emerging. To quote Kohane [42], “Clinical decision support algorithms will be derived entirely from data, not expert opinion, market incentives, or committee consensus. The huge amount of data available will make it possible to draw inferences from observations that will not be encumbered by unknown confounding”. Being able to base statistical methodological research for early phase dose finding on observed health care data is now possible.

In this context, the objective would be to identify novel dose-finding approaches taking into account cumulative toxicities over cycles using dynamic treatment regime methodology. Additionally, instead of estimating a single dose as the MTD, the aim could be to recommend a dose regimen. This would require detailed information concerning the everyday care of the patients in hospitals (dose modification, acute and cumulative toxicities, ...). Real clinical data would allow to evaluate the “actions” undertaken for subsequent “observations” and if there are “regimen” patterns or subgroups that are given repeatedly or if there are as many dose-regimens as patients. Then, novel designs incorporating patterns in dose-regimen allocation could be developed. The patient care data cannot be analyzed using standard regression model approaches. Each observation will induce an action (stopping temporarily or definitively the treatment, decreasing, reducing or maintaining the dose) by the physicians (each physician according to his own experience will not necessarily undertake the same action regarding the same observation) which can accrue several times over the entire cycles of treatment. In this case, observations and actions would be considered random variables. Dynamic treatment regime approaches, in which treatment choices made for a particular patient are based on that individual’s characteristics and history with the goal of optimizing his or her long-term clinical outcome, can be used to analyze data and design new dose-finding methods [1, 89, 32, 80].

Another important perspective in the context of phase I dose-finding clinical trials is the incorporation of pharmacokinetic/pharmacodynamic (PK/PD) data in the dose determination process. Indeed, until now the PK/PD analysis and the dose-finding scheme are completely separated. Experts including physicians, statisticians, and pharmacologists are meeting after each cohort inclusion to discuss and share their analysis in order to determine the next dose level to administer. Nevertheless, both PK/PD and statistical dose-finding analysis are not theoretically modeled together in order to take into account all available data. Therefore, a new challenge is to combine these two analyses in the dose selection process. This subject is, in my opinion, particularly relevant and interesting.

A new Bayesian dose-finding method for an ordinal efficacy endpoint and a binary toxicity endpoint

Most phase I/II approaches consider the efficacy response as a binary endpoint (success or failure). However, in oncology the efficacy response, such as the tumoral size decreases, is usually given in terms of regression (RE), stability (ST), partial response (PR) and complete response (CR) using RECIST (Response Evaluation Criteria In Solid Tumors). Therefore, modelling efficacy as an ordinal endpoint seems relevant and correspond to clinical reality. Moreover, dichotomising efficacy responses will also reduce all available information. We would like to propose a new dose-finding method for phase I/II studies where the toxic response is modelled as a binary variable and the efficacy response is modeled as an ordinal variable.

We would like to use a simple 2-parameter dose-free logistic model for toxicity with parameters constrained such as the toxicity is increasing with the dose level. For efficacy, a 6-parameter multinomial logistic regression model would be used with parameters constrained to ensure that the probability of RE is decreasing with the dose.

A utility function will be constructed to select the best dose level. Depending of the pathology, different weights could be used on the efficacy responses: regression, stability, partial response and complete response.

This project is still on-going and is developed in collaboration with Dr. Monia Ezzalfani.

Phase I/II Dose-Finding Design for Molecularly Targeted Agent with Continuous Efficacy Outcome

We would like to extend our paper “Phase I/II Dose-Finding Design for Molecularly Targeted Agent: Plateau Determination using Adaptive Randomization” to continuous efficacy variable between 0 and 1. Indeed, after discussions with statisticians and physicians inside IRIS, it appeared that in early phases the efficacy outcome is often considered as a percentage of reduction of some molecules or percentage of achieving the target. Therefore, for use in current practice our aim is to propose an extension of our design. We would also like to study the impact of patient variability in efficacy response.

A simple 2-parameter logistic model would be considered for toxicity. As previously, we assumed that efficacy monotonically increases with the dose and then plateau after reaching the level of saturation. Let ϕ_k denote the efficacy probability for dose level k . The efficacy outcome, z , is continuous between 0 and 1. Linear regression model are not appropriate for situations where the continuous response is restricted to the interval $[0;1]$ since it may yield fitted values that exceed this interval. As proposed by Ferrari and Cribari-Neto [24], we used a simple beta regression model appropriate for the case where the efficacy is measured continuously on the standard unit interval.

The beta distribution is flexible to model proportions since its density can

have many different shapes depending on the two parameters of its distribution. After a re-parametrization of the beta distribution in terms of efficacy mean and precision parameter, the efficacy mean can be modeled using a simple logistic model (see appendix A.5). All toxicity and efficacy parameters will be estimated by MCMC, as well as posterior toxicity and efficacy probability distributions. In the simulation study, we would like to simulate patients' efficacy responses using a normal distribution with a variance more or less small in order to study the impact of patients' variability in efficacy responses on the performance of the design.

For the plateau determination, the same criterion could be used with posterior probabilities and adaptive randomization. Nevertheless, we would like to try another criterion, the posterior predictive loss, that we tried but did not work in the case of binary outcome (see Appendix A.6). This research line should be explored to see what would be applicable in this context or how to adapt it.

Appendix A

Appendices

A.1 Lee and Cheung's skeleton

Lee and Cheung studied the influence of the initial guesses in order to find the bests to optimize PCS.

Calculate indifference intervals from the initial guesses

Let $\pi(d, a)$ be the supposed dose-toxicity relationship. We assumed that we have L dose levels (d_1, \dots, d_L) and the target probability toxicity is θ . Let $\mathcal{A} = [A_1, A_{L+1}]$ be the parameter space (i.e. $a \in \mathcal{A}$), and

$$\begin{cases} H_1 = [A_1, A_2] \\ H_\ell = [A_\ell, A_{\ell+1}] \quad \forall \ell \in \llbracket 2, L-1 \rrbracket \\ H_L = [A_L, A_{L+1}] \end{cases}$$

where A_ℓ is the solution for $\pi(d_{\ell-1}, A_\ell) + \pi(d_\ell, A_\ell) = 2\theta \quad \forall \ell \in \llbracket 2, L \rrbracket$. It means that with this parameter A_ℓ , the estimated mean of probabilities of toxicity of doses d_ℓ and $d_{\ell-1}$ is equal to the target.

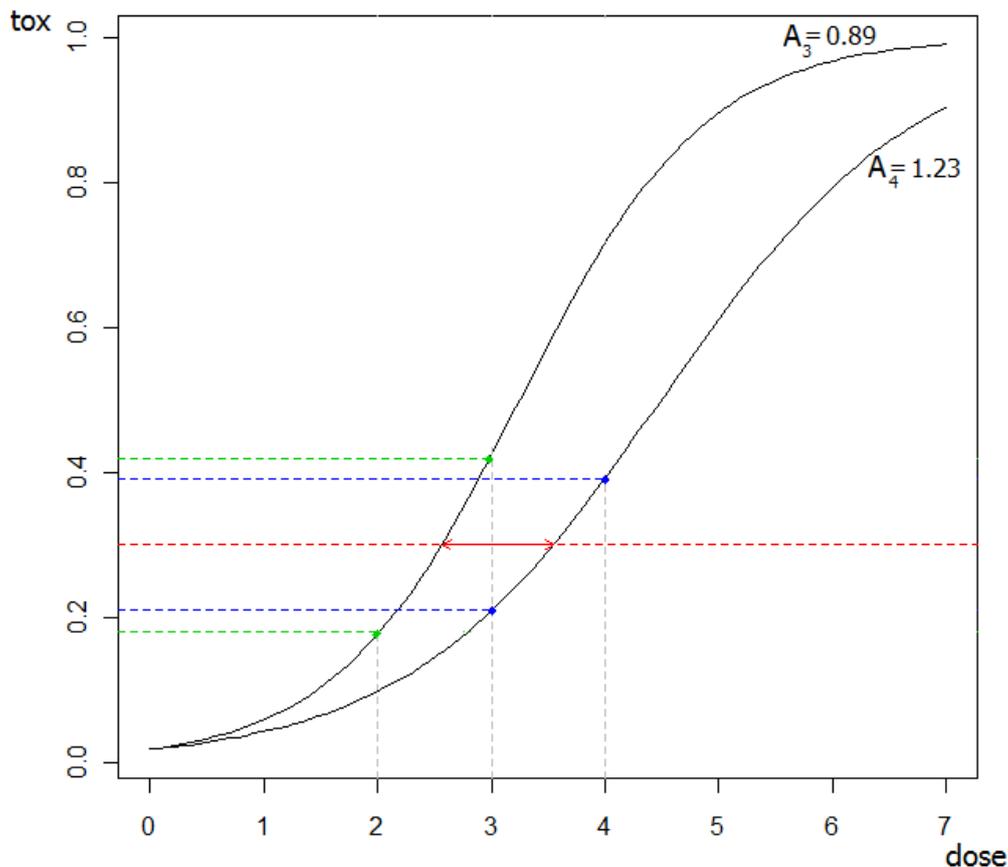


Figure A.1: Illustration of the construction of H intervals.

To understand the principle, if we take the example of Figure A.1. For $\ell = 3$, we are looking for A_3 such as the mean between probabilities of toxicity of dose 2 and 3, given by the model with this parameter value, is equal to the target. In the example, we had $A_3 = 0.89$. Similarly, for $\ell = 4$, we are looking for A_4 such as the mean between probabilities of toxicity of dose 3 and 4, given by the model with this parameter value, is equal to the target. We had $A_4 = 1.23$. Therefore $H_3 = [0.89, 1.23]$.

Shen and O'Quigley showed that for large enough n the CRM will recommend the true MTD (m) with certainty, if $a_\ell \in H_m \forall \ell$, where a_ℓ is defined such that $\pi(d_\ell, a_\ell) = \mu_\ell$ and μ_ℓ is the true toxicity probability associated with dose ℓ . In the example, if every $a_k \in [0.89, 1.23]$ then dose 3 will be the MTD, which is intuitive on Figure A.1.

Cheung and Chappell postulated that if the true dose toxicity function is step around the MTD, for large enough n , the dose recommended by the CRM is the true MTD under the more relaxed conditions whereby:

$$\begin{cases} a_m \in H_m \\ a_\ell \in \bigcup_{i=\ell+1}^L H_i \quad \forall \ell \in \llbracket 1, m-1 \rrbracket \\ a_\ell \in \bigcup_{i=1}^{\ell-1} H_i \quad \forall \ell \in \llbracket m+1, L \rrbracket \end{cases}$$

In fact, the practical use is very limited because the conditions involve the unknown true probabilities of toxicity μ_ℓ 's. Therefore, Cheung and Chappell suggested converting the intervals in the parameter space for a into intervals on the probability of toxicity scale $\pi_\ell = \pi(d_\ell, a)$. Then, the indifference interval for a given correct dose level m was defined as an interval of probabilities of toxicity associated with the neighbouring doses such that these neighbouring doses may be selected instead of the true MTD (m). The indifference for the MTD (m) will be denoted by

$$\begin{cases} [NA, \pi(d_{m+1}, A_{m+1})] \text{ for } m = 1 \\ [\pi(d_{m-1}, A_m), \pi(d_{M+1}, A_{M+1})] \quad \forall m \in \llbracket 2, L-1 \rrbracket \\ [\pi(d_{m-1}, A_m), NA] \text{ for } m = L \end{cases}$$

Let us take an example. Assume that $\theta = 0.25$, and the dose-toxicity relationship is supposed to be empiric, i.e. $\pi(d_\ell, a) = w_\ell^a$ where $a \in]0, 5]$, and the initial guesses of probabilities of toxicity at each of the 5 doses are $w = (0.05, 0.12, 0.25, 0.40, 0.55)$.

$$\begin{aligned} A_2 \text{ is defined such that: } d_1^{A_2} + d_2^{A_2} &= 2\theta = 0.5 \\ \Rightarrow A_2 &= 0.553. \end{aligned}$$

Similarly, all A_ℓ could be calculated and then the sets are: $H_1 =]0, 0.553]$, $H_2 = [0.553, 0.816]$, $H_3 = [0.816, 1.241]$, $H_4 = [1.241, 1.89]$ and $H_5 = [1.89, 5]$.

If $m = 3$, the condition specified by Cheung and Chappell is:

$$a_3 \in H_3, a_1 \in \bigcup_{i=2}^5 H_i, a_2 \in \bigcup_{i=3}^5 H_i, a_4 \in \bigcup_{i=1}^3 H_i \text{ and } a_5 \in \bigcup_{i=1}^4 H_i.$$

Then because $\pi(d_\ell, a_\ell) = \mu_\ell$, the intervals of μ_ℓ can be obtained by:

$$[\mu_{\ell,inf}, \mu_{\ell,sup}] = [\pi(d_\ell, a_{\ell,sup}), \pi(d_\ell, a_{\ell,inf})].$$

For example $\mu_3 \in [0.25^{1.241}, 0.25^{0.816}] = [0.179, 0.323]$.

We obtained: $\mu_1 \in [0, 0.191]$, $\mu_2 \in [0, 0.177]$, $\mu_3 \in [0.179, 0.323]$, $\mu_4 \in [0.321, 1]$, and

$\mu_5 \in [0.323, 1]$. For n large enough, with the chosen working model, if the true probabilities of toxicity are in these intervals, the true MTD will be selected.

If $\mu_2 \in [0.177, \mu_3]$, the CRM can select dose 2 as the MTD instead of dose 3, but the probability of toxicity could be consider close enough to be indifferent to select dose 2 instead of dose 3. It is the same if $\mu_4 \in [\mu_3, 0.321]$. So the indifference interval in the case where $m = 3$ is $[0.177, 0.321]$. The indifference intervals can be calculated assuming that each dose is the MTD. Then we obtain $[NA, 0.31]$, $[0.19, 0.32]$, $[0.18, 0.32]$, $[0.18, 0.32]$, $[0.18, NA]$. (Of note: The overall indifference interval is the union of these intervals, in this case $[0.18, 0.32]$. Therefore when the target toxicity and the working model are specified, it can be known when CRM could failed.)

Calculate initial guesses from indifference intervals

The idea here is to do the contrary, i.e. when the target toxicity θ , the prior MTD m , the number of dose L and the length of the indifference interval δ are specified, using a model $\pi(d_\ell, a)$, the working model can be obtained using backward substitution in the precedent method.

Indeed, with the empiric working model $\pi(d_\ell, a) = w_\ell^a$, at the beginning $w_m = \theta$ (because $\pi(d_\ell, \hat{a}_0) = \theta$ where \hat{a}_0 is the prior mean of a).

Indifference intervals of length 2δ then w_{m-1} and w_{m+1} can be obtained using:

$$\begin{cases} \pi(d_{m-1}, A_m) + \pi(d_m, A_m) = 2\theta \\ \text{and} \\ \pi(d_m, A_{m+1}) + \pi(d_{m+1}, A_{m+1}) = 2\theta \end{cases} \quad \text{which is based on the definition of } A_\ell.$$

and

$$\begin{cases} \pi(d_{m-1}, A_m) = \theta - \delta \\ \text{and} \\ \pi(d_{m+1}, A_{m+1}) = \theta + \delta \end{cases} \quad \begin{array}{l} \text{based on the definition of indifference interval} \\ \text{(the indifference interval for the MTD } m \text{ is} \\ [\pi(d_{m-1}, A_m), \pi(d_{m+1}, A_{m+1})]) \text{ and its length } \delta. \end{array}$$

If we decide that the overall indifference interval is the same than indifference intervals for each dose, we can similarly use the same procedure to obtain iteratively w_1, \dots, w_L .

Lee and Cheung showed that as δ increases, the distance between w_ℓ ($\ell \neq m$) and θ also increases.

They made extensive simulations to find, given a number of doses L , a prior MTD m , a sample size N and a target toxicity θ , the best δ to optimize the percentage of correct selection (PCS).

The global idea

To summarize the global idea of Lee and Cheung's skeleton:

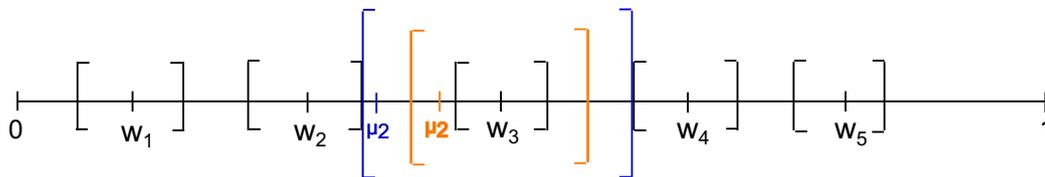


Figure A.2: The idea of Lee and Cheung's skeleton.

On Figure A.2, w are the initial guesses, and according to Cheung and Chappell, if the true probabilities of toxicity are in the black intervals, the true MTD will be selected. If we assume that the true MTD is dose 3, the indifference interval for dose 3 is the blue one. It means that if the true probabilities of the doses next to dose 3 are in this interval, they could be selected instead of dose 3. For example on Figure A.2, μ_2 , the true probability of toxicity of dose 2, is in this interval so dose 2 could be selected instead of dose 3. Now we do it in the opposite direction. We choose the length of indifference interval, so we can take the one we want around dose 3, for example the interval in orange. Therefore we can choose how much the other doses should be close to be selected instead of the good one. But we can't hope to have small indifference interval and initial guesses very spaced, because the smaller indifference interval is, the closer initial guesses are. Therefore, we have to find the best length of indifference intervals to improve PCS. That what Lee and Cheung did next in their article.

A.2 Clopper-Pearson Confidence Interval

Confidence interval based on normal approximation

To determine a confidence interval for a proportion, we know the formula based on normal approximation when $np \geq 5$ and $n(1 - p) \geq 5$:

$$\hat{p} \pm z_{1-\frac{\alpha}{2}} \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

where p is the proportion of interest, n is the sample size, α is the desired confidence and $z_{1-\frac{\alpha}{2}}$ is the quantile of order $1 - \frac{\alpha}{2}$ of a normal distribution.

Clopper-Pearson confidence interval

When $np < 5$ and $n(1 - p) < 5$ or when $p = 0$ or $p = 1$, the previous formula does not work. Clopper and Pearson developed a method to construct confidence interval in [14].

The Clopper-Pearson interval is an early and very common method for calculating binomial confidence intervals.[3] This is often called an 'exact' method, but that is because it is based on the cumulative probabilities of the binomial distribution (i.e. exactly the correct distribution rather than an approximation), but the intervals are not exact in the way that one might assume: the discontinuous nature of the binomial distribution precludes any interval with exact coverage for all population proportions. The Clopper-Pearson interval can be written as

$$\left\{ p \mid P(\mathcal{B}(n, p) \leq X) \geq \frac{\alpha}{2} \right\} \cap \left\{ p \mid P(\mathcal{B}(n, p) \geq X) \geq \frac{\alpha}{2} \right\} \quad (\text{A.1})$$

where X is the number of successes observed in the sample and $\mathcal{B}(n, p)$ is a binomial random variable with n trials and probability of success p .

This interval never has less than the nominal coverage for any population proportion, but that means that it is usually conservative. To calculate Clopper-Pearson interval the function "*binom.test*" from basic package or "*binconf*" from *Hmisc* library can be used.

Construction of Clopper-Pearson interval

The idea is that normally, we are looking for (p_{\min}, p_{\max}) such as:

$$P(p_{\min} < p < p_{\max}) = 1 - \alpha \quad (\text{A.2})$$

With the normal approximation, this equation leads to equation (A.2). When it could not be used, we have to find another method. I saw two different ways to present you how to obtain Clopper-Pearson interval.

The first is that when we are doing a binomial test, we are testing:

$(H_0)p = p_0$: the parameter p of the binomial distribution is equal to p_0 vs.
 $(H_1)p \neq p_0$: the parameter p of the binomial distribution is different from p_0 ,
 with an α risk.

It means that we have $1 - \alpha$ % of chance to accept $(H_0)p = p_0$ if it is true. Therefore a confidence interval for p at $1 - \alpha$ would be all the values for which the test would be accepted (not rejected).

The two-sided test statistic for the exact binomial test consists first in finding c_1 and c_2 such as:

$$\sum_{i=0}^{c_1} P(\mathcal{B}(n, p_0) = i) = P(\mathcal{B}(n, p_0) \leq c_1) \leq \frac{\alpha}{2}$$

and

$$\sum_{i=c_2}^n P(\mathcal{B}(n, p_0) = i) = P(\mathcal{B}(n, p_0) \geq c_2) \leq \frac{\alpha}{2}$$

Then (H_0) is accepted if the number of observed events $x \in]c_1, c_2[$ and rejected otherwise.

So here, as we said, we want to find all the values for p such as $x \in]c_1, c_2[$. According to the test statistics, x will be in $]c_1, c_2[$ if:

$$P(\mathcal{B}(n, p_0) \leq x) \geq \frac{\alpha}{2}$$

and

$$P(\mathcal{B}(n, p_0) \geq x) \geq \frac{\alpha}{2}$$

Therefore, a confidence interval for p is:

$$\left\{ p \mid P(\mathcal{B}(n, p) \leq X) \geq \frac{\alpha}{2} \right\} \cap \left\{ p \mid P(\mathcal{B}(n, p) \geq X) \geq \frac{\alpha}{2} \right\}$$

and so the exact Clopper-Pearson confidence interval inverts two single-tailed Binomial test at the desired alpha.

Another way to see it is that instead of searching directly solving equation (A.2), we want to find all the values of parameter p such as what happened had at least a probability of α to happen. (To determine a confidence interval, we are looking for all the values which could lead to this result of events, except the α less probable values). Therefore we will choose to have at most a probability of $1 - \frac{\alpha}{2}$ on both sides not it to happen. If the event is " $\mathcal{B}(n; p) = X$ ", then the confidence interval is:

$$\left\{ p \mid P(\mathcal{B}(n, p) < X) \leq 1 - \frac{\alpha}{2} \right\} \cap \left\{ p \mid P(\mathcal{B}(n, p) > X) \leq 1 - \frac{\alpha}{2} \right\}$$

Rewritten, it corresponds to Clopper-Pearson interval given in equation (A.1).

A.3 Adaptive Rejection Metropolis Sampling within Gibbs Sampling and reminders

Rejection Sampling

The rejection sampling is a method for drawing independent samples from a distribution f . If we have a density g from which we can draw samples easily and such as it exists a constant c such that $\forall x \in \mathcal{D}, cg(x) \geq f(x)$, where \mathcal{D} is the domain of f , then we can simulate a variable whose density is f .

For k from 0 to L

- (1) Sample x_k from g
- (2) Sample u from Uniform[0, 1]
- (3) If $u > \frac{f(x_k)}{cg(x_k)}$ then rejection : return to (1)
Else acceptance: $y_k = x_k$

\Rightarrow return (y_0, \dots, y_L)

The idea is that: we have c and g such as $0 \leq \frac{f(x)}{cg(x)} \leq 1$.

So when we simulate $u \sim \mathbb{U}[0, 1]$, we accept $X \sim g$ with probability $\frac{f(x)}{cg(x)}$.

The closest f is from cg , the closest the probability $\frac{f(x)}{cg(x)}$ is from 1, so $X \sim f$.

The expected number of iterations is c because the number of iterations required follows a geometric law of parameter $P(\text{accept})$.

Yet,

$$P(X = x \text{ is accepted}) = P(X=x) \times P(\text{accept}|X = x)$$

Therefore,

$$\begin{aligned} P(\text{accept}) &= \int_{\mathcal{D}} g(x) \times \frac{f(x)}{cg(x)} dx \\ &= \int_{\mathcal{D}} \frac{f(x)}{c} dx \\ &= \frac{1}{c} \end{aligned}$$

The number of iterations $\sim \mathcal{Geometric}(\frac{1}{c})$, and so the expected number of iterations is c .

Adaptive Rejection Sampling

The idea of the adaptive rejection sampling is to reduce the number of iterations by improving the envelope function g after each rejection. After rejection, we want to include the new information we have about f , which is assumed to be log-concave.

We say that a function f is log-concave if:

$$\forall a, b, c \in \mathcal{D} / a < b < c, \ln(f(a)) - 2\ln(f(b)) + \ln(f(c)) < 0$$

This definition does not assume continuity in derivatives of f .

Let $S_n = \{x_i, i \in \llbracket 0, n+1 \rrbracket\}$ be a current set of abscissae in ascending order (x_0 and x_{n+1} can be infinite lower and upper limits of \mathcal{D}).

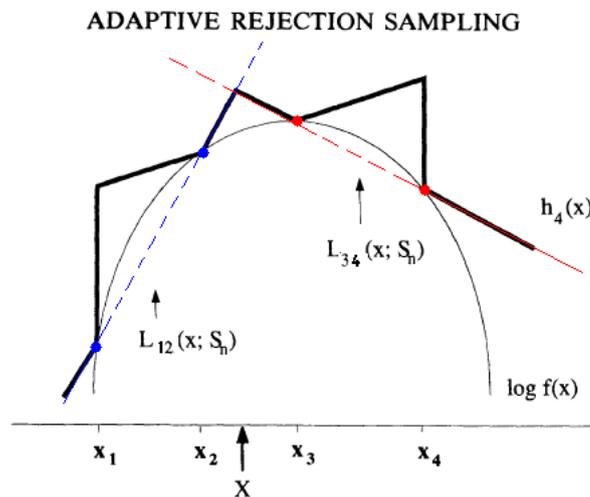
$\forall 1 \leq i \leq j \leq n$, let $L_{i,j}(x, S_n)$ be the straight line through points $(x_i, \ln(f(x_i)))$ and $(x_j, \ln(f(x_j)))$ and for other (i, j) , $L_{i,j}(x, S_n)$ is undefined.

The envelop is defined such as:

$$\forall x / x_i \leq x < x_{i+1}, h_n(x) = \min(L_{i-1,i}(x, S_n), L_{i+1,i+2}(x, S_n))$$

By convention, if b is undefined then $\min(a, b) = \min(b, a) = a$.

Because f is log-concave, h_n is an envelope for $\ln(f)$, i.e. $\forall x \in \mathcal{D}, h_n(x) \geq \ln(f(x))$.



Adaptive rejection function $h_4(x)$ (—) for a log-concave function $f(x)$: X is sampled from $h_4(x)$

Figure A.3: Construction of the function envelop for ARS.

Indeed, because f is log-concave, if we take two points on f , the segment joining these two points will be below f . So if we extend the segment into a line, the other parts will be above f .

So, we have $h_n(x) \geq \ln(f(x))$

$\Leftrightarrow \exp(h_n(x)) \geq f(x)$ because \exp is strictly increasing on \mathbb{R}

To be a density, let

$$g_n(x) = \frac{1}{c_n} \exp(h_n(x))$$

where

$$c_n = \int_{\mathcal{D}} \exp(h_n(x)) dx$$

And because g_n is a piecewise exponential, it can be sampled directly.

The ARS algorithm is now the following:

For k from 0 to L

(1) Initialize n and S_n

(2) Sample x_k from g

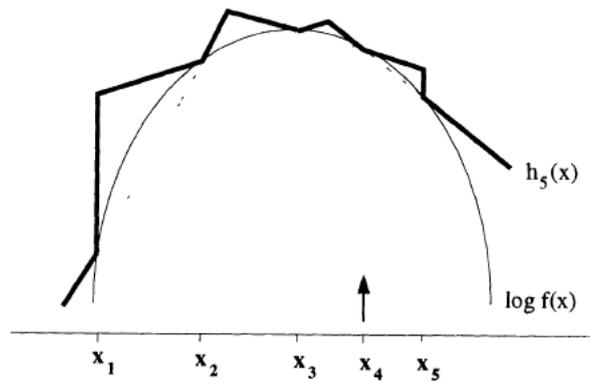
(3) Sample u from Uniform[0, 1]

(4) If $u > \frac{f(x_k)}{\exp(h_n(x_k))}$ then rejection : $S_{n+1} = S_n \cup \{x_k\}$, reorder in ascending order S_{n+1} , increment n and return to (2)

Else acceptance: $y_k = x_k$

\Rightarrow return (y_0, \dots, y_L)

We can notice that after each rejection, the number of abscissae points is increasing (and so the number of contact points with $\ln(f)$), so c is decreasing together with the probability of rejection.



Updated current set S_5 and rejection function $h_5(x)$ (—) after incorporating X

Figure A.4: Evolution of the envelope after several rejection.

If \mathcal{D} is not bounded on the left, x_0 should be chosen such as the gradient of $L_{1,2}(x, S_n)$ is positive, and if \mathcal{D} is not bounded on the right, x_{n+1} should be chosen such as the gradient of $L_{n-1,n}(x, S_n)$ is negative.

If f is not log-concave, this algorithm can not be used, because g_n is not necessarily an envelope of f .

Hastings-Metropolis algorithm

The Metropolis algorithm is an MCMC (Markov Chains Monte Carlo) method. This method is used to simulate a law f when classical methods are not efficient and/or when the density is known except for a normalisation constant.

We choose a proposal distribution $q(\cdot|\cdot)$ defined as a transition kernel or as a conditional law. Then the algorithm is the following:

- (1) Initialize y_0
- For k from 0 to L
- (2) Sample x_k from $q(\cdot|y_k)$
 - (3) Sample u from Uniform[0, 1]
 - (4) If $u > \min\left(1, \frac{f(x_k)q(y_k|x_k)}{f(y_k)q(x_k|y_k)}\right)$ then rejection : $y_{k+1} = y_k$
 Else acceptance: $y_{k+1} = x_k$
- \Rightarrow return (y_0, \dots, y_L)

This Markov Chain is convergent according to ergodic theorem to the law of f .

Tierney (1991) suggested the use of the Hastings-Metropolis algorithm within Gibbs sampling to sample from full conditional distributions. For this, y_0 should be the value of y at the start of the current Gibbs iteration, and y_1 will be the new value for y . But the chain may be slower to converge.

Gibbs Sampling

Gibbs sampling is an algorithm to generate a sequence of samples from the joint probability distribution of two or more random variables. The purpose of such a sequence is to approximate the joint distribution, to approximate the marginal distribution of one of the variables, or some subset of the variables, or to compute an integral.

The Gibbs sampler relies on the availability of all complete conditional distributions. Let $\theta = (\theta_1, \dots, \theta_n)$ the vector of the parameters of the model.

The idea is that we start from an arbitrary point $\theta^{(0)} = (\theta_1^{(0)}, \dots, \theta_n^{(0)}) \in \Theta$, we sampled in turn in each of the full conditional distribution in updating them as the process goes on.

Let $\theta_{(-i)} = (\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_n)$ the vector θ without its i^{th} component. The density of the full conditional distribution of θ_i is $f(\theta_i | \theta_{(-i)}, \text{data})$. Let us assume that we know its closed form for all $i \in \{1, \dots, n\}$.

If we want to sample from a joint distribution, the algorithm can be described as follow:

(1) Initialize $\theta^{(0)} = (\theta_1^{(0)}, \dots, \theta_n^{(0)})$

For m from 1 to $M + N$

(2) Sample

- $\theta_1^{(m)}$ from $f(\theta_1 | \theta_2^{(m-1)}, \dots, \theta_n^{(m-1)}, \text{data})$
- \dots
- $\theta_i^{(m)}$ from $f(\theta_i | \theta_1^{(m)}, \dots, \theta_{i-1}^{(m)}, \theta_{i+1}^{(m-1)}, \dots, \theta_n^{(m-1)}, \text{data})$
- \dots
- $\theta_n^{(m)}$ from $f(\theta_n | \theta_1^{(m)}, \dots, \theta_{n-1}^{(m)}, \text{data})$

\Rightarrow return $(\theta^{(M+1)}, \dots, \theta^{(M+N)})$

M is a number of iterations that will be discarded. They are called "burn-in" iterations which correspond to the time before convergence.

It can be shown that $\theta^{(M)} = (\theta_1^{(M)}, \dots, \theta_n^{(M)})$ converges in distribution to the posterior joint distribution $f(\theta_1, \dots, \theta_n | \text{data})$. The following N values are then considered as a sample from this distribution.

Adaptive Rejection Metropolis Sampling within Gibbs sampling

ARS can not be used to sample from non log-concave distributions. When this is the case, Gilks, Best and Tan propose to replace the rejection sampling in favour of the Hastings-Metropolis algorithm to update one parameter at a time. But to avoid high probabilities of rejection, they adapted the proposal density q to the shape of the full conditional density f using ARS. They added to ARS a single Hastings-Metropolis step thus creating ARMS within Gibbs chain. (ARMS will not produce independent samples from f unlike ARS).

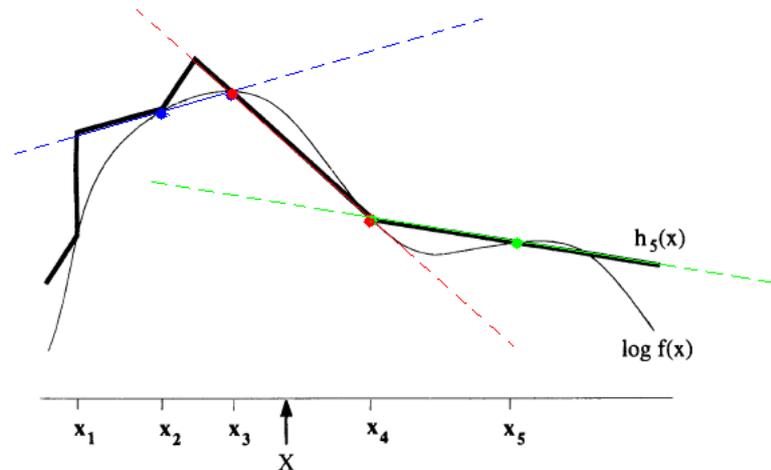
Let (Y, Z) denote the complete set of variables being sampled by the Gibbs sampler. Y is the current variable to be sampled from its full conditional density f . Let y_{cur} denote the current value of x at a given iteration of the Gibbs sampler. The aim then is to replace y_{cur} with a new value y from f .

h is now defined as follow:

$$\forall x / x_i \leq x \leq x_{i+1}, h_n(x) = \max(L_{i,i+1}(x, S_n), \min(L_{i-1,i}(x, S_n), L_{i+1,i+2}(x, S_n)))$$

By convention, if b is undefined then $\min(a, b) = \min(b, a) = \max(a, b) = \max(b, a) = a$.

In general, h_n won't be an envelope of $\ln(f)$:



Adaptive rejection function $h_5(x)$ (—) for a non-log-concave function $f(x)$

Figure A.5: Adaptive rejection function for ARMS.

When the shape of $\log(f)$ is concave, the same lines than in ARS are used to envelop $\log(f)$ because h_n reduces to its previous expression, and when the shape of $\log(f)$ becomes convex, it is the straight line which is used. So in general, h_n won't be an envelope of $\log(f)$.

g_n is defined as previously:

$$g_n(x) = \frac{1}{c_n} \exp(h_n(x))$$

where

$$c_n = \int_{\mathcal{D}} \exp(h_n(x)) dx$$

Starting abscissa for ARMS must be independent of y_{cur} . The algorithm can be described in the following way:

- (1) Initialize n and S_n independently of y_{cur}
 - (2) Sample x_k from g_n
 - (3) Sample u from Uniform $[0, 1]$
 - (4) If $u > \frac{f(x_k)}{\exp(h_n(x_k))}$ then ARS rejection: $S_{n+1} = S_n \cup \{x_k\}$, reorder in ascending order S_{n+1} , increment n and return to (2)
Else ARS acceptance: $y_{\text{temp}} = x_k$
 - (5) Sample u from Uniform $[0, 1]$
 - (6) If $u > \min \left(1, \frac{f(y_{\text{temp}}) \min(f(y_{\text{cur}}), \exp(h_n(y_{\text{cur}})))}{f(y_{\text{cur}}) \min(f(y_{\text{temp}}), \exp(h_n(y_{\text{temp}})))} \right)$ then Metropolis rejection : $y = y_{\text{cur}}$
Else Metropolis acceptance: $y = y_{\text{temp}}$
- \Rightarrow return y

As we said, if f is log-concave, ARMS reduces to ARS because h_n will be an envelope for f , so $\min(f(x), \exp(h_n(x))) = f(x) \forall x$ and so (6) will be always accepted.

The proof that ARS preserves the stationary distribution of the Gibbs sampler can be found in [28].

A.4 Isotonic regression and Pool-Adjacent-Violators Algorithm

Isotonic Regression

Isotonic regression consists in projecting a non-parametric function in the set of monotonic non-decreasing functions.

1-dimensional isotonic regression

A function g on \mathcal{G} is isotonic if it is increasing (non-decreasing) on \mathcal{G} .

Definition: Let g be a function on \mathcal{G} . g^* on \mathcal{G} is an isotonic regression of g with weights w if and only if g^* is isotonic and minimize:

$$\sum_{x \in \mathcal{G}} w_x (g(x) - f(x))^2$$

on the set of isotonic functions on \mathcal{G} .

Several algorithms exist to implement isotonic regression, the most common is PAVA which stand for "Pooled Adjacent Violator Algorithm". The idea of the algorithm consist in replacing the two variables where the increasing constraint is not respected by their weighted mean and to repeat this process while the set is not increasing. This algorithm is detailed below:

- If $g(x)$ is non-decreasing, i.e. $g(x) \leq g(x + 1)$ then let:

$$g^*(x) = g(x)$$

- Else, somewhere the increasing constraint is not respected, i.e.

$$\exists x_v / g(x) > g(x_v + 1)$$

The two values are replaced by their weighted mean:

$$\frac{w_{x_v} g(x_v) + w_{x_v+1} g(x_v + 1)}{w_{x_v} + w_{x_v+1}}$$

Then the elements $(x_v, x_v + 1)$ form a "block"

- If this new set of values (of $\#\mathcal{G} - 1$ elements) is isotonic, then let:

$$g^*(x_v) = g^*(x_v + 1) = \frac{w_{x_v} g(x_v) + w_{x_v+1} g(x_v + 1)}{w_{x_v} + w_{x_v+1}} \quad \text{for } x_v \text{ and } x_v + 1,$$

and $g^*(x) = g(x)$ for other elements

- If this new set is not isotonic, we repeat the process on it

Example: $g(\mathcal{G}) = \{0, 2, 4, 3\}$

\Rightarrow constraint not respected between the 3rd et 4th element

$w = \{1, 1, 3, 1\}$

then: $g^*(\mathcal{G}) = \left\{ 0, 2, \frac{4 \times 3 + 3 \times 1}{3 + 1}, \frac{4 \times 3 + 3 \times 1}{3 + 1} \right\} = \{0, 2, 3.75, 3.75\}$

The PAVA algorithm is already implemented on R via the function "*pava*" in package *Iso*.

2-dimensional isotonic regression

In the 2-dimensional isotonic regression, the function must be non-decreasing on each of its two parameters.

Let $\Omega = \{(i, j), i = 1, \dots, I, j = 1, \dots, J\}$. Let g a function on this set, let G denote its matrix: $G = (g_{i,j}) = (g(i, j))$.

$F : \Omega \mapsto R$ is isotonic if $\forall i, j, k, \ell / (i, j) \leq (k, \ell)$ (i.e. $i \leq k$ and $j \leq \ell$), $f_{i,j} \leq f_{k,\ell}$. It means that F is non-decreasing along lines and columns.

G^* on Ω is an isotonic regression of G with weights w if and only if G^* is isotonic and minimize:

$$\sum_{(i,j) \in \Omega} w_{i,j} (g_{i,j} - f_{i,j})^2$$

on the set of isotonic functions on Ω .

An algorithm was proposed by Dykstra and Robertson in the 2-dimensional case. This is based on an algorithm in the 1-dimensional case and is detailed below:

- Let $\hat{G}^{(1)} = (\hat{g}_{i,j}^{(1)})$ denote the solution of the isotonic regression of $G = (g_{i,j})$ on lines, i.e. $\hat{G}^{(1)}$ minimize

$$\sum_{i=1}^I w_{i,j} (g_{i,j} - f_{i,j})^2 \text{ with } f_{1,j} \leq \dots \leq f_{I,j} \text{ for } j = 1, \dots, J.$$

$$\text{Let } R^{(1)} = (r_{i,j}^{(1)}) = (\hat{g}_{i,j}^{(1)} - g_{i,j})$$

- Let $\tilde{G}^{(1)} = (\tilde{g}_{i,j}^{(1)})$ denote the solution of the isotonic regression of $G + R^{(1)}$ on columns, i.e. $\tilde{G}^{(1)}$ minimize

$$\sum_{j=1}^J w_{i,j} (g_{i,j} + r_{i,j}^{(1)} - f_{i,j})^2 \text{ with } f_{i,1} \leq \dots \leq f_{i,J} \text{ for } i = 1, \dots, I.$$

$$\text{Let } C^{(1)} = \tilde{G}^{(1)} - (G + R^{(1)})$$

- In the n -th iteration:

$\hat{G}^{(n)}$ is obtained by performing the isotonic regression of $G + C^{(n-1)}$ on the lines.

$$\text{Let } R^{(n)} = \hat{G}^{(n)} - (G + C^{(n-1)}).$$

$$\text{Therefore } \hat{G}^{(n)} = G + C^{(n-1)} + R^{(n)}.$$

Then $\tilde{G}^{(n)}$ is obtained by performing the isotonic regression of $G + R^{(n)}$ on columns.

$$\text{Let } C^{(n)} = \tilde{G}^{(n)} - (G + R^{(n)}).$$

$$\text{Therefore } \tilde{G}^{(n)} = G + R^{(n)} + C^{(n)}.$$

Dykstra and Robertson showed that $\tilde{G}^{(n)}$ and $\hat{G}^{(n)}$ converge both to G^* as $n \rightarrow +\infty$.

An algorithm enabling to apply an isotonic regression is available on R via the "biviso" function in package *Iso* (provided that the given matrix does not contain missing values).

A.5 Beta regression

Linear regression models are not appropriate for situations where the continuous response is restricted to the interval $[0;1]$ since it may yield fitted values that exceed this interval. Ferrari and Cribari-Neto [24] proposed a simple beta regression model appropriate for the case where the efficacy is measured continuously on the standard unit interval. Indeed, the beta distribution is flexible to model proportions since its density can have many different shapes depending on the two parameters of its distribution (Figure A.6).

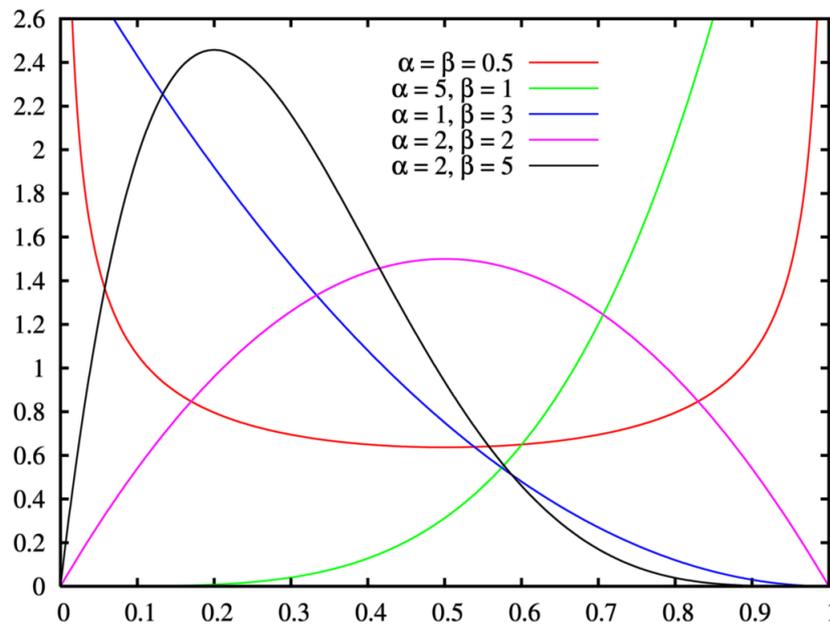


Figure A.6: Various shapes of beta densities.

The beta density is given by:

$$\pi(z; p, q) = \frac{\Gamma(p+q)}{\Gamma(p)\Gamma(q)} z^{p-1} (1-z)^{q-1}$$

where $0 < z < 1$, $p > 0$, $q > 0$ and Γ is the gamma function. The mean and variance of the beta distribution is given by:

$$E(z) = \frac{p}{p+q}$$

and

$$V(z) = \frac{pq}{(p+q)^2(p+q+1)}$$

The beta density can be re-parametrized in terms of parameters (μ, ϕ) instead of (p, q) such as:

$$\begin{cases} \mu = \frac{p}{p+q} \\ \phi = p+q \end{cases} \Leftrightarrow \begin{cases} p = \mu\phi \\ q = (1-\mu)\phi \end{cases}$$

μ is the mean of the response variable z and ϕ can be interpreted as a precision parameter, as for fixed μ , the larger the value of ϕ , the smaller the variance of z . It follows the mean, variance, and re-parametrized density:

$$E(z) = \mu$$

$$V(z) = \frac{\mu(1-\mu)}{1+\phi} = \frac{V(\mu)}{1+\phi}$$

and

$$\pi(z; \mu, \phi) = \frac{\Gamma(\phi)}{\Gamma(\mu\phi)\Gamma((1-\mu)\phi)} z^{\mu\phi-1} (1-z)^{(1-\mu)\phi-1}$$

with $0 < z < 1$, $0 < \mu < 1$, and $\phi > 0$.

Finally, the mean proportion can be modeled using different link functions as for instance a logit function:

$$\text{logit}(\mu_k) = \gamma_0 + \gamma_1 v_k$$

A.6 Posterior Predictive Loss

The Posterior Predictive Loss (PPL) is a model selection criterion proposed by Gelfand and Ghosh [27]. The PPL is a criterion which minimize the expectation of prediction error and is defined by:

$$D_k(\tau_v) = \sum_{\ell=1}^I \min_{a_l} (E_{z_{\ell,\text{rep}}|\mathcal{D}_{\text{eff}},\tau_v} (\text{loss}(z_{\ell,\text{rep}}, a_l)) + k \cdot \text{loss}(z_{\ell,\text{obs}}, a_l))$$

where $z_{\ell,\text{rep}}$ replicate of $z_{\ell,\text{obs}}$ and k compromise between the goodness-of-fit of the model and a penalty term .

Several conditions on the *loss* function are required: (1) the existence of the second partial derivatives (loss_{02} and loss_{20}), (2) the loss function must be non-negative, and (3) the loss function must satisfy $\text{loss}(b, b) = 0$ and $\text{loss}_{01}(b, b) = 0$. A simple and well-known loss function is the squared error loss: $\text{loss}(y, a) = (y - a)^2$. This choice of function is convenient as it enables for $D_k(\tau_v)$ to be explicitable:

$$D_k(\tau_v) = \sum_{\ell=1}^I V(z_{\ell,\text{rep}}|\mathcal{D}_{\text{eff}},\tau_v) + \frac{k}{k+1} \sum_{\ell=1}^I (E(z_{\ell,\text{rep}}|\mathcal{D}_{\text{eff}},\tau_v) - z_{\ell,\text{obs}})^2$$

where the expectation and the variance are with respect to the posterior predictive distribution associated with $z_{\ell,\text{rep}}$ under τ_v .

Other general classes of loss functions were proposed under certain conditions where Taylor's series are used to approximate D_k .

Bibliography

- [1] ALMIRALL, D., COMPTON, S. N., GUNLICKS-STOESSEL, M., DUAN, N., AND MURPHY, S. A. Designing a pilot sequential multiple assignment randomized trial for developing an adaptive treatment strategy. *Stat Med* 31, 17 (Jul 2012), 1887–1902.
- [2] ARROWSMITH, J. Trial watch: phase III and submission failures: 2007-2010. *Nat Rev Drug Discov* 10, 2 (Feb 2011), 87.
- [3] ASCIERTO, P. A., AND MARINCOLA, F. M. Combination therapy: the next opportunity and challenge of medicine. *J Transl Med* 9 (2011), 115.
- [4] AZRIEL, D., MANDEL, M., AND RINOTT, Y. The treatment versus experimentation dilemma in dose-finding studies. *Journal of Statistical Planning and Inference* 141 (2011), 2759–2768.
- [5] BABB, J., ROGATKO, A., AND ZACKS, S. Cancer phase I clinical trials: efficient dose escalation with overdose control. *Stat Med* 17, 10 (May 1998), 1103–1120.
- [6] BAILEY, S., NEUENSCHWANDER, B., LAIRD, G., AND BRANSON, M. A Bayesian case study in oncology Phase I combination dose-finding using logistic regression with covariates. *J Biopharm Stat* 19 (2009), 469–484.
- [7] BEKELE, B. N., AND SHEN, Y. A Bayesian approach to jointly modeling toxicity and biomarker expression in a phase I/II dose-finding trial. *Biometrics* 61, 2 (Jun 2005), 343–354.
- [8] BRAUN, T. M., AND WANG, S. A hierarchical Bayesian design for phase I trials of novel combinations of cancer therapeutic agents. *Biometrics* 66, 3 (Sep 2010), 805–812.
- [9] CHEN, Z., TIGHIOUART, M., AND KOWALSKI, J. Dose escalation with overdose control using a quasi-continuous toxicity score in cancer Phase I clinical trials. *Contemp Clin Trials* 33, 5 (Sep 2012), 949–958.
- [10] CHEUNG, Y. K., AND CHAPPELL, R. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* 56, 4 (Dec 2000), 1177–1182.

- [11] CHEUNG, Y. K., AND CHAPPELL, R. A simple technique to evaluate model sensitivity in the continual reassessment method. *Biometrics* 58 (2002), 671–674.
- [12] CHEVRET, S. The continual reassessment method in cancer phase I clinical trials: a simulation study. *Stat Med* 12, 12 (Jun 1993), 1093–1108.
- [13] CHEVRET, S. *Statistical Methods for dose-Finding Experiments*. Statistics in Practice. John Wiley and Sons Ltd., Chichester, 2006.
- [14] CLOPPER, C. J., AND PEARSON, E. S. The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial. *Biometrika* 26 (1934), 404–13.
- [15] CONAWAY, M. R., DUNBAR, S., AND PEDDADA, S. D. Designs for single- or multiple-agent phase I trials. *Biometrics* 60 (Sep 2004), 661–669.
- [16] CROWLEY, J., AND PAULER ANKERST, D. *Handbook of statistics in clinical oncology*. Chapman and Hall/CRC, Boca Raton, 2006.
- [17] DAIMON, T., ZOHAR, S., AND O’QUIGLEY, J. Posterior maximization and averaging for Bayesian working model choice in the continual reassessment method. *Stat Med* 30, 13 (Jun 2011), 1563–1573.
- [18] DOUSSAU, A., THIEBAUT, R., AND PAOLETTI, X. Dose-finding design using mixed-effect proportional odds model for longitudinal graded toxicity data in phase I oncology clinical trials. *Stat Med* 32, 30 (Dec 2013), 5430–5447.
- [19] DRUKER, B. J. Perspectives on the development of a molecularly targeted agent. *Cancer Cell* 1, 1 (Feb 2002), 31–36.
- [20] DURHAM, S. D., FLOURNOY, N., AND ROSENBERGER, W. F. A random walk rule for phase I clinical trials. *Biometrics* 53, 2 (Jun 1997), 745–760.
- [21] DYKSTRA, R. L., AND ROBERTSON, T. An algorithm for isotonic regression for two or more independent variables. *Ann Stat* 10 (1982), 708–716.
- [22] FAN, S. K., VENOOK, A. P., AND LU, Y. Design issues in dose-finding Phase I trials for combinations of two agents. *J Biopharm Stat* 19 (2009), 509–523.
- [23] FARIES, D. Practical modifications of the continual reassessment method for phase I cancer clinical trials. *J Biopharm Stat* 4, 2 (Jul 1994), 147–164.
- [24] FERRARI, S., AND CRIBARI-NETO, F. Beta regression for modelling rates and proportions. *Journal of Applied Statistics* 31, 7 (2004), 799–815.
- [25] GAFTER-GVILI, A., LEADER, A., GURION, R., VIDAL, L., RAM, R., SHACHAM-ABULAFIA, A., BEN-BASSAT, I., LISHNER, M., SHPILBERG, O., AND RAANANI, P. High-dose imatinib for newly diagnosed chronic phase

- chronic myeloid leukemia patients—systematic review and meta-analysis. *Am. J. Hematol.* 86, 8 (Aug 2011), 657–662.
- [26] GARRETT-MAYER, E. The continual reassessment method for dose-finding studies: a tutorial. *Clin Trials* 3, 1 (2006), 57–71.
- [27] GELFAND, A., AND S., G. Model choice: a minimum posterior predictive loss approach. *Biometrika* 85 (1998), 1–11.
- [28] GILKS, W. R., BEST, N. G., AND TAN, K. K. C. Adaptive Rejection Metropolis Sampling within Gibbs Sampling. *JRSS* 44 (1995), 455–472.
- [29] GITTINS, J. Bandit processes and dynamic allocation indices. *JRSS, Series B* 41 (1979), 148–177.
- [30] GOODMAN, S. N., ZAHURAK, M. L., AND PIANTADOSI, S. Some practical improvements in the continual reassessment method for phase I studies. *Stat Med* 14, 11 (Jun 1995), 1149–1161.
- [31] HAINES, L. M., PEREVOZSKAYA, I., AND ROSENBERGER, W. F. Bayesian optimal designs for Phase I clinical trials. *Biometrics* 59, 3 (Sep 2003), 591–600.
- [32] HERNAN, M. A., LANOY, E., COSTAGLIOLA, D., AND ROBINS, J. M. Comparison of dynamic treatment regimes via inverse probability weighting. *Basic Clin. Pharmacol. Toxicol.* 98, 3 (Mar 2006), 237–242.
- [33] HOERING, A., LEBLANC, M., AND CROWLEY, J. Seamless phase I-II trial design for assessing toxicity and efficacy for targeted agents. *Clin. Cancer Res.* 17, 4 (Feb 2011), 640–646.
- [34] HRYNIUK, W. M. More is better. *J. Clin. Oncol.* 6, 9 (Sep 1988), 1365–1367.
- [35] HUANG, X., BISWAS, S., OKI, Y., ISSA, J. P., AND BERRY, D. A. A parallel phase I/II clinical trial design for combination therapies. *Biometrics* 63, 2 (Jun 2007), 429–436.
- [36] HUNSBERGER, S., RUBINSTEIN, L. V., DANCEY, J., AND KORN, E. L. Dose escalation trial designs based on a molecularly targeted endpoint. *Stat Med* 24, 14 (Jul 2005), 2171–2181.
- [37] IASONOS, A., WILTON, A. S., RIEDEL, E. R., SESHAN, V. E., AND SPRIGGS, D. R. A comprehensive comparison of the continual reassessment method to the standard 3+3 dose escalation scheme in Phase I dose-finding studies. *Clin Trials* 5, 5 (2008), 465–477.
- [38] IVANOVA, A. Escalation, group and A + B designs for dose-finding trials. *Stat Med* 25, 21 (Nov 2006), 3668–3678.

- [39] IVANOVA, A., MONTAZER-HAGHIGHI, A., MOHANTY, S. G., AND DURHAM, S. D. Improved up-and-down designs for phase I trials. *Stat Med* 22, 1 (Jan 2003), 69–82.
- [40] IVANOVA, A., AND WANG, K. A non-parametric approach to the design and analysis of two-dimensional dose-finding trials. *Stat Med* 23 (2004), 1861–1870.
- [41] IVANOVA, A., AND XIAO, C. Dose finding when the target dose is on a plateau of a dose-response curve: comparison of fully sequential designs. *Pharm Stat* 12, 5 (Sep 2013), 309–314.
- [42] KOHANE, I. S., DRAZEN, J. M., AND CAMPION, E. W. A glimpse of the next 100 years in medicine. *N. Engl. J. Med.* 367, 26 (Dec 2012), 2538–2539.
- [43] KORN, E. L., MIDTHUNE, D., CHEN, T. T., RUBINSTEIN, L. V., CHRISTIAN, M. C., AND SIMON, R. M. A comparison of two phase I trial designs. *Stat Med* 13 (1994), 1799–806.
- [44] LE TOURNEAU, C., DIERAS, V., TRESCA, P., CACHEUX, W., AND PAOLETTI, X. Current challenges for the early clinical development of anticancer drugs in the era of molecularly targeted agents. *Target Oncol* 5, 1 (Mar 2010), 65–72.
- [45] LE TOURNEAU, C., GAN, H. K., RAZAK, A. R., AND PAOLETTI, X. Efficiency of new dose escalation designs in dose-finding phase I trials of molecularly targeted agents. *PLoS ONE* 7, 12 (2012), e51039.
- [46] LE TOURNEAU, C., RAZAK, A. R., GAN, H. K., POP, S., DIERAS, V., TRESCA, P., AND PAOLETTI, X. Heterogeneity in the definition of dose-limiting toxicity in phase I cancer clinical trials of molecularly targeted agents: a review of the literature. *Eur. J. Cancer* 47, 10 (Jul 2011), 1468–1475.
- [47] LEE, S., AND CHEUNG, Y. K. Model calibration in the continual reassessment method. *Clin Trials* 6 (2009), 227–38.
- [48] LEE, S. M., AND CHEUNG, Y. K. Calibration of prior variance in the Bayesian continual reassessment method. *Stat Med* 30, 17 (Jul 2011), 2081–2089.
- [49] LEVY, V., ZOHAR, S., BARDIN, C., VEKHOFF, A., CHAOU, D., RIO, B., LEGRAND, O., SENTENAC, S., ROUSSELOT, P., RAFFOUX, E., CHAST, F., CHEVRET, S., AND MARIE, J. P. A phase I dose-finding and pharmacokinetic study of subcutaneous semisynthetic homoharringtonine (ssHHT) in patients with advanced acute myeloid leukaemia. *Br. J. Cancer* 95, 3 (Aug 2006), 253–259.

- [50] MANDREKAR, S. J., CUI, Y., AND SARGENT, D. J. An adaptive phase I design for identifying a biologically optimal dose for dual agent drug combinations. *Stat Med* 26 (May 2007), 2317–2330.
- [51] MANDREKAR, S. J., QIN, R., AND SARGENT, D. J. Model-based phase I designs incorporating toxicity and efficacy for single and dual agent drug combinations: methods and challenges. *Stat Med* 29 (May 2010), 1077–1083.
- [52] MAUGUEN, A., LE DELEY, M. C., AND ZOHAR, S. Dose-finding approach for dose escalation with overdose control considering incomplete observations. *Stat Med* 30, 13 (Jun 2011), 1584–1594.
- [53] NEUENSCHWANDER, B., BRANSON, M., AND GSPONER, T. Critical aspects of the Bayesian approach to phase I cancer trials. *Stat Med* 27, 13 (Jun 2008), 2420–2439.
- [54] O’QUIGLEY, J., PEPE, M., AND FISHER, L. Continual reassessment method: a practical design for phase 1 clinical trials in cancer. *Biometrics* 46 (1990), 33–48.
- [55] O’QUIGLEY, J., AND REINER, E. A stopping rule for the continual reassessment method. *Biometrika* 85, 3 (1998), 741–748.
- [56] O’QUIGLEY, J., AND SHEN, L. Z. Continual reassessment method: a likelihood approach. *Biometrics* 52, 2 (Jun 1996), 673–684.
- [57] O’QUIGLEY, J., AND ZOHAR, S. Experimental designs for phase I and phase I/II dose-finding studies. *Br. J. Cancer* 94, 5 (Mar 2006), 609–613.
- [58] O’QUIGLEY, J., AND ZOHAR, S. Retrospective robustness of the continual reassessment method. *J Biopharm Stat* 20, 5 (Sep 2010), 1013–1025.
- [59] ORON, A. P., AND HOFF, P. D. Small-sample behavior of novel phase I cancer trial designs. *Clin Trials* 10, 1 (Feb 2013), 63–80.
- [60] PAOLETTI, X., AND KRAMAR, A. A comparison of model choices for the Continual Reassessment Method in phase I cancer trials. *Stat Med* 28, 24 (Oct 2009), 3012–3028.
- [61] POLLEY, M. Y. Practical modifications to the time-to-event continual reassessment method for phase I cancer trials with fast patient accrual and late-onset toxicities. *Stat Med* 30, 17 (Jul 2011), 2130–2143.
- [62] POSTEL-VINAY, S., ARKENAU, H. T., OLMOS, D., ANG, J., BARRIUSO, J., ASHLEY, S., BANERJI, U., DE-BONO, J., JUDSON, I., AND KAYE, S. Clinical benefit in Phase-I trials of novel molecularly targeted agents: does dose matter? *Br. J. Cancer* 100, 9 (May 2009), 1373–1378.

- [63] ROBBINS, H. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society* 58 (1952), 527–535.
- [64] ROGATKO, A., SCHOENECK, D., JONAS, W., TIGHIOUART, M., KHURI, F. R., AND PORTER, A. Translation of innovative designs into phase I trials. *J. Clin. Oncol.* 25, 31 (Nov 2007), 4982–4986.
- [65] ROSENBERGER, W. F., AND HAINES, L. M. Competing designs for phase I clinical trials: a review. *Stat Med* 21, 18 (Sep 2002), 2757–2770.
- [66] SIMON, R., FREIDLIN, B., RUBINSTEIN, L., ARBUCK, S. G., COLLINS, J., AND CHRISTIAN, M. C. Accelerated titration designs for phase I clinical trials in oncology. *J. Natl. Cancer Inst.* 89, 15 (Aug 1997), 1138–1147.
- [67] STORER, B. E. Design and analysis of phase I clinical trials. *Biometrics* 45 (1989), 925–37.
- [68] SUTTON, R., AND BARTO, A. Reinforcement learning: An introduction. *Cambridge, MA: MIT Press* (1998).
- [69] THALL, P. F., AND COOK, J. D. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* 60, 3 (Sep 2004), 684–693.
- [70] THALL, P. F., MILLIKAN, R. E., MUELLER, P., AND LEE, S. J. Dose-finding with two agents in Phase I oncology trials. *Biometrics* 59, 3 (Sep 2003), 487–496.
- [71] THALL, P. F., AND NGUYEN, H. Q. Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *J Biopharm Stat* 22, 4 (2012), 785–801.
- [72] THALL, P. F., AND RUSSELL, K. E. A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* 54, 1 (Mar 1998), 251–264.
- [73] TIGHIOUART, M., LIU, Y., AND ROGATKO, A. Escalation with overdose control using time to toxicity for cancer phase I clinical trials. *PLoS ONE* 9, 3 (2014), e93070.
- [74] TIGHIOUART, M., AND ROGATKO, A. Dose finding with escalation with overdose control (ewoc) in cancer clinical trials. *Statistical Science* 25, 2 (05 2010), 217–226.
- [75] TING, N. *Dose Finding in Drug Development*. Statistics for Biology and Health. Springer, 2006.
- [76] U.S. DEPARTMENT OF HEALTH, HUMAN SERVICES, F., DRUG ADMINISTRATION, C. F. B. E., AND RESEARCH. Guidance for Industry, Clinical Considerations for Therapeutic Cancer Vaccines.

- [77] WAGES, N. A., CONAWAY, M. R., AND O'QUIGLEY, J. Continual Reassessment Method for Partial Ordering. *Biometrics* 67, 4 (Dec 2011), 1555–1563.
- [78] WAGES, N. A., CONAWAY, M. R., AND O'QUIGLEY, J. Dose-finding design for multi-drug combinations. *Clin Trials* 8 (2011), 380–389.
- [79] WANG, K., AND IVANOVA, A. Two-dimensional dose finding in discrete dose space. *Biometrics* 61 (2005), 217–222.
- [80] WANG, L., ROTNITZKY, A., LIN, X., MILLIKAN, R., AND THALL, P. F. Evaluation of viable dynamic treatment regimes in a sequentially randomized trial of advanced prostate cancer. *Journal of the American Statistical Association* 107, 498 (2012), 493–508.
- [81] WHITEHEAD, J. Bayesian decision procedures with application to dose-finding studies. *International Journal of Pharmaceutical Medicine* 11 (1997), 201–208.
- [82] WHITEHEAD, J., THYGESEN, H., AND WHITEHEAD, A. A Bayesian dose-finding procedure for phase I clinical trials based only on the assumption of monotonicity. *Stat Med* 29 (Jul 2010), 1808–1824.
- [83] WHITEHEAD, J., THYGESEN, H., AND WHITEHEAD, A. Bayesian procedures for phase I/II clinical trials investigating the safety and efficacy of drug combinations. *Stat Med* 30 (Jul 2011), 1952–1970.
- [84] YIN, G., LI, Y., AND JI, Y. Bayesian dose-finding in phase I/II clinical trials using toxicity and efficacy odds ratios. *Biometrics* 62, 3 (Sep 2006), 777–784.
- [85] YIN, G., AND YUAN, Y. A latent contingency table approach to dose finding for combinations of two agents. *Biometrics* 65 (2009), 866–875.
- [86] YIN, G., AND YUAN, Y. Bayesian dose finding in oncology for drug combinations by copula regression. *JRSS* 58 (2009), 211–224.
- [87] YUAN, Y., AND YIN, G. Sequential continual reassessment method for two-dimensional dose finding. *Stat Med* 27, 27 (Nov 2008), 5664–5678.
- [88] YUAN, Y., AND YIN, G. Robust EM Continual Reassessment Method in Oncology Dose Finding. *Journal of the American Statistical Association* 106 (Nov 2011), 818–831.
- [89] ZHANG, B., TSIATIS, A. A., LABER, E. B., AND DAVIDIAN, M. Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika* 100, 3 (2013).
- [90] ZHANG, J., BRAUN, T. M., AND TAYLOR, J. M. Adaptive prior variance calibration in the Bayesian continual reassessment method. *Stat Med* 32, 13 (Jun 2013), 2221–2234.

- [91] ZHANG, W., SARGENT, D. J., AND MANDREKAR, S. An adaptive dose-finding design incorporating both toxicity and efficacy. *Stat Med* 25 (Jul 2006), 2365–2383.
- [92] ZHOU, Y. Choice of designs and doses for early phase trials. *Fundam Clin Pharmacol* 18, 3 (Jun 2004), 373–378.
- [93] ZOHAR, S., AND CHEVRET, S. The continual reassessment method: comparison of Bayesian stopping rules for dose-ranging studies. *Stat Med* 20, 19 (Oct 2001), 2827–2843.
- [94] ZOHAR, S., AND LEVY, V. Dose Estimation: A Key Step in Malignancies Drug Development. *Pharm Med* 22 (2008), 35–40.

Résumé long

Les essais de phases I sont généralement les premiers essais cliniques destinés à tester l'administration d'un nouveau traitement chez l'être humain. Bien que ce nouveau médicament ait été intensivement expérimenté sur des animaux de laboratoire, les effets indésirables chez l'Homme ne peuvent pas toujours être anticipés. Ainsi, l'objectif des essais de phase I est d'évaluer la toxicité du nouveau traitement et d'en identifier ses effets indésirables. En cancérologie, les phases I en sont conduites sur des malades plutôt que sur des volontaires sains, à cause de la forte nocivité des traitements anti-cancéreux dont l'utilisation chez ces derniers ne serait pas éthique. De plus, bien que ces études n'aient pas pour but de déterminer l'efficacité du traitement contre le cancer, il y a quand même un fort intérêt, en tant que critère secondaire de l'étude, à pouvoir observer l'efficacité potentielle du nouveau traitement chez ces patients. Les patients inclus lors de cette phase précoce sont souvent à un stade avancé de la maladie et plusieurs lignes de traitements ont déjà été testées et d'avèrent avoir échoué. Le nouvel agent développé est parfois leur dernière option thérapeutique. Les essais de phase I en cancérologie ne recrutent qu'un faible nombre de participants, généralement entre 15 et 50. Pour les agents cytotoxiques, qui représentent la majorité des agents testés jusqu'à nos jours, il est supposé que (1) la toxicité augmente avec la dose administrée, et que (2) l'efficacité est nécessairement positivement corrélée avec la toxicité. Ainsi, plus la dose est élevée, plus elle est supposée être toxique, mais également efficace. Ceci définit le paradigme « Plus est mieux ». L'objectif principal des essais de cliniques de phase I en cancérologie est donc de déterminer la dose la plus forte du nouveau traitement qui peut être administrée, tout en conservant un taux de toxicité acceptable. Cette dose est appelée dose maximale tolérée. La problématique statistique est de sélectionner, parmi un ensemble de doses prédéfinies, la dose dont la probabilité de toxicité est la plus proche d'une toxicité cible fixée a priori (Figure A.7).

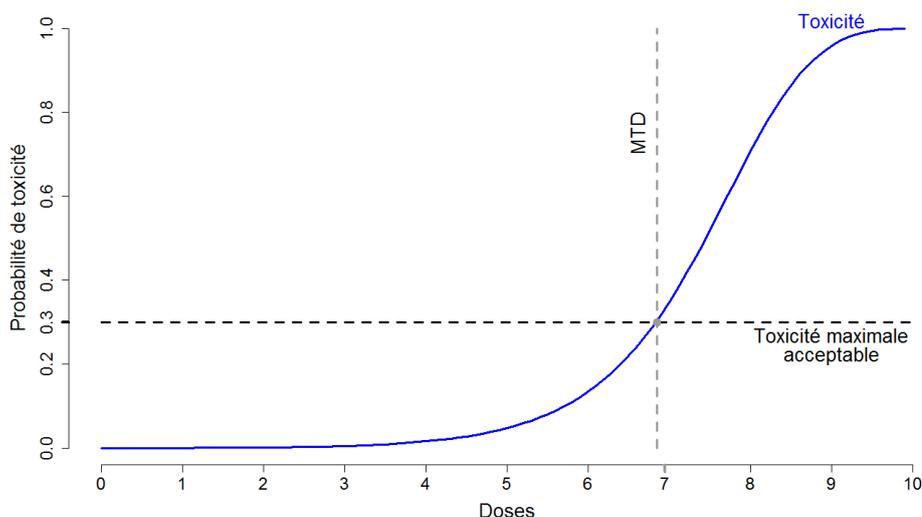


Figure A.7: Détermination de la dose maximale tolérée.

Ces essais de phases précoces sont des procédures séquentielles d'escalade de doses. Cela consiste à graduellement escalader (ou dé-escalader) les niveaux de dose de l'agent testé jusqu'à ce que la toxicité observée liée au médicament atteigne un seuil prédéfini maximum acceptable.

Au cours de ces dernières années, les oncologistes ont prescrits des chimiothérapies standards car ils ont pu observer et prouver que ces agents étaient efficaces. Ces agents permettent de faire régresser le cancer chez certains patients en tuant les cellules à division rapide. Cependant, elles entraînent souvent des dommages collatéraux sur les tissus sains, créant ainsi des effets indésirables qui altèrent de nombreux systèmes dont le système circulatoire, le système immunitaire, et le système digestif. En effet, comme les agents cytotoxiques perturbent habituellement les molécules et réactions chimiques qui se produisent dans les cellules à division rapide, de nombreuses cellules normales, situées dans tout le corps humain, qui subissent une croissance active et sont en division cellulaire peuvent aussi être endommagées par la chimiothérapie. La plupart des designs statistiques pour les essais de phase I se trouvant dans la littérature, comme le « 3+3 », la méthode de réévaluation séquentielle ou l'escalade de dose avec contrôle du surdosage, ont été développés selon les hypothèses liées aux agents cytotoxiques. Mais, depuis quelques années, des thérapies ciblées comme « les petites molécules » ou « les anticorps » ont vu le jour. Contrairement aux chimiothérapies standards, les thérapies ciblées sont conçues pour interagir avec des molécules spécifiques intervenant dans des processus utilisés par les cellules cancéreuses pour croître, se diviser et se propager dans tout le corps. Lorsque les chercheurs découvrent une molécule potentiellement vulnérable impliquée dans un processus du développement du cancer, ils conçoivent de nouvelles thérapies visant à interrompre son

activité. Pour les thérapies ciblées, les effets secondaires sont souvent moindres que pour la chimiothérapie standard, car elles causent peu voire pas de dommages collatéraux aux cellules saines. Cela peut contribuer à améliorer la qualité de vie des patients suivant ces traitements.

Dans ce contexte, nous avons décidé de développer une méthode de recherche de dose en monothérapie pour thérapie ciblée. Ce design peut être utilisé pour les molécules ciblées en combinaison avec une dose fixe d'agent cytotoxique, ce qui est en voie d'extension dans les essais cliniques sur le cancer. Ce travail a été initié à la fois suite à une nécessité pratique lors d'un véritable essai clinique au sein de l'institut de recherche internationales Servier, laboratoire pharmaceutique avec lequel j'ai accompli ma thèse, mais également suite à des discussions avec des médecins des hôpitaux publiques concernant la différence d'hypothèses entre les thérapies ciblées et les cytotoxiques. Après plusieurs discussions avec des médecins et pharmacologues, il en est ressorti que l'efficacité était supposée croissante avec la dose puis atteignait un plateau (d'efficacité). En effet, lorsque tous les récepteurs ciblés sont déjà liés au nouveau traitement, il n'est pas nécessaire d'augmenter le niveau de dose puisque la saturation du corps s'avère être atteinte. Pour notre étude, l'efficacité était supposée être un critère binaire. Nous avons donc proposé un design de recherche de dose bayésien de phase I/II pour les thérapies ciblées. Pour modéliser la relation dose-toxicité et dose-efficacité; nous avons utilisé des modèles logistiques en intégrant une 1-spline dans le modèle d'efficacité afin de modéliser le plateau. Notre méthode se concentre sur la sélection de la dose optimale, c'est-à-dire sur la dose associée à la toxicité la plus faible parmi celles ayant la plus forte efficacité, plutôt que sur la dose maximale tolérée, ceci dans le but de réduire la toxicité pour la même efficacité (Figure A.8).

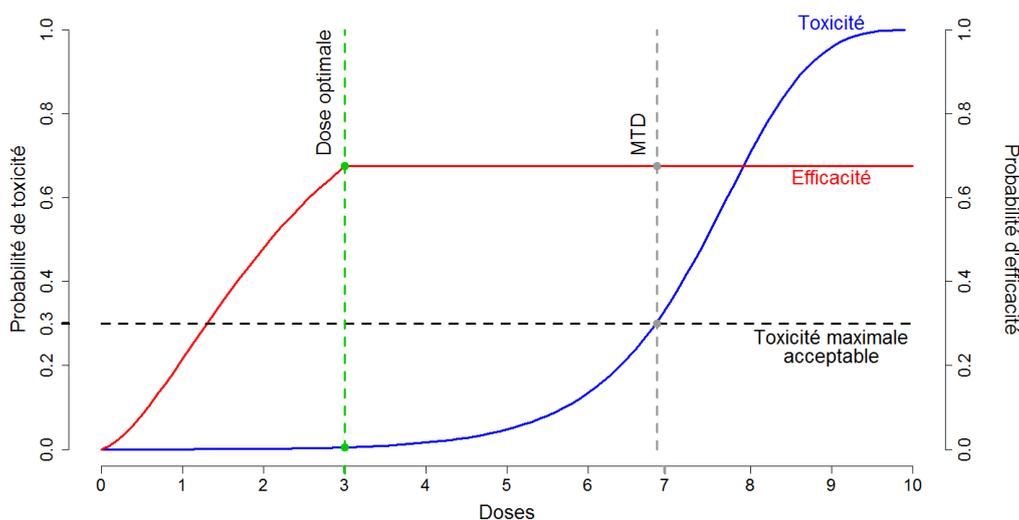


Figure A.8: Détermination de la dose optimale.

Lors de l'élaboration de cette méthode adaptative, nous avons rencontré des problèmes dans l'estimation du plateau. En effet, il a été reconnu dans la prise de décisions séquentielles que les algorithmes qui choisissent chaque action successive en optimisant un critère de décision peuvent se retrouver coincés sur une action sous-optimale. Cela est dû au fait que l'algorithme sélectionne de façon répétée l'action sous-optimale et ne permet donc pas d'accumuler suffisamment de données pour sélectionner la véritable action optimale. Ce problème est parfois connu comme le dilemme « optimisation versus exploration » et a été reconnu dans le cadre des essais cliniques de recherche de dose. Pour éviter ce problème, nous avons utilisé la randomisation adaptative dans notre processus d'allocation de doses. Celle-ci permet de tirer aléatoirement le paramètre d'intérêt dans l'ensemble de ses valeurs possibles selon les probabilités a posteriori estimées. Par conséquent, il permet d'utiliser les informations accumulées à travers les probabilités utilisées, mais aussi d'ajouter une part de hasard nécessaire afin de débloquer l'algorithme. Les deux méthodes d'allocation proposées donnent de bons résultats, comparables en termes de pourcentage de sélection correcte de la dose optimale. Cependant, la méthode basée sur les probabilités a posteriori de la position du plateau semble plus robuste sur les différents scénarios testés. En effet, les pourcentages de sélection correcte sont toujours supérieurs à 50%. En outre, elle donne également de meilleurs résultats en termes de pourcentage de sélection d'une dose correcte, c'est à dire d'une dose ayant la plus forte efficacité, mais parmi elles pas nécessairement la plus faible toxicité. Malheureusement, pour notre essai clinique, en raison de l'hétérogénéité des patients à inclure pour cette phase I, l'efficacité ne pouvait pas être évaluée avec précision, et donc ne pouvait pas être prise en compte dans la conception du design. La méthode proposée a finalement été abandonnée et remplacée par une CRM. Néanmoins, le développement de thérapies ciblées devenant pratique courante, un autre essai clinique à l'intérieur de l'entreprise pourrait prochainement être mis en place en utilisant ce design. Il s'agira d'une combinaison d'une molécule ciblée avec une dose fixe de chimiothérapie standard. Nous avons également étendu notre design (1) aux relations unimodales, et (2) à différents groupes d'efficacité selon un biomarqueur, menant à la recommandation d'une dose optimale différente dans chaque sous-groupe, mais partageant une toxicité commune. Nous avons observé qu'en général, la méthode proposée semble avoir de bonnes performances. Pour des questions de temps, cette méthode a été implémentée en C/C++ qui est beaucoup plus rapide que R. Pour être facilement utilisée dans la pratique, nous sommes actuellement en train de développer un package R. Ce package permettra à la fois d'effectuer des simulations d'essais cliniques avec un choix flexible de paramètres en entrée, et également d'estimer la prochaine dose à administrer à partir des données d'un véritable essai clinique. Nous prévoyons de terminer l'implémentation de ce package pour la soutenance de thèse.

Avec les progrès récents dans le domaine de l'oncologie, il est de plus en plus rare de trouver de nouvelles molécules plus performantes que les stratégies thérapeutiques existantes. En outre, les cancers peuvent développer divers mécanismes

de résistance à un traitement impliquant un seul agent. C'est pourquoi, dans différents domaines, mais surtout dans les études sur le cancer, de plus en plus d'études de combinaisons sont mises en place. En combinant plusieurs agents, les investigateurs espèrent augmenter l'action anti-tumorale et la survie globale des patients grâce un effet de synergie entre les agents en termes d'efficacité. Cependant, lorsque l'on combine plusieurs agents, l'ordre des combinaisons en termes de probabilité de toxicité n'est pas complètement connu. Par exemple, la combinaison de deux agents cytotoxiques, pour lesquels la toxicité est croissante avec chaque agent n'induit qu'une relation d'ordre partielle. En effet, si l'on fixe un agent, la toxicité de la combinaison augmente lorsque l'on augmente la dose de l'autre agent.

$$\begin{array}{cccccc}
 D_{1,3} & < & D_{2,3} & < & D_{3,3} & < & D_{4,3} & < & D_{5,3} \\
 \mathbf{v} & & \mathbf{v} & & \mathbf{v} & & \mathbf{v} & & \mathbf{v} \\
 D_{1,2} & < & D_{2,2} & < & D_{3,2} & < & D_{4,2} & < & D_{5,2} \\
 \mathbf{v} & & \mathbf{v} & & \mathbf{v} & & \mathbf{v} & & \mathbf{v} \\
 D_{1,1} & < & D_{2,1} & < & D_{3,1} & < & D_{4,1} & < & D_{5,1}
 \end{array}$$

Figure A.9: Relation d'ordre partielle entre les combinaisons de deux agents cytotoxiques.

Même si un ordre partiel est connu, il est tout de même difficile de choisir comment augmenter ou diminuer une combinaison de doses. En effet, sur une diagonale, il n'y a aucune connaissance de la combinaison la plus toxique. Par exemple, sur la figure A.9, on ne sait pas laquelle entre la combinaison $D_{1,2}$ et $D_{2,1}$ est la plus toxique a priori. Par conséquent, il n'est pas raisonnable d'utiliser des méthodes de recherche de dose de monothérapies pour des études de combinaisons.

Nous avons effectué une revue systématique de la littérature de tous les essais cliniques de phase I de combinaisons de molécules publiés les trois dernières années entre le 1^{er} Janvier 2011 et le 31 Décembre 2013, où le niveau de dose d'au moins deux agents varie. Notre objectif était de déterminer quelles étaient les pratiques actuelles des essais de combinaisons. Notre analyse a mis en évidence que tous les designs utilisés étaient destinés à la monothérapie. Ainsi la relation dose-toxicité était ramenée à un espace de dose unidimensionnel alors que la réalité impliquait plusieurs agents induisant un problème multi-dimensionnel. En particulier, la plupart des essais (88%) ont utilisé le traditionnel « 3+3 » ou une modification de ce dernier. Pour ramener le problème à un espace à une dimension, les médecins ont présélectionné les combinaisons à évaluer en choisissant un ensemble de combinaisons dont l'ordre de toxicité est connu. En effet, 62.7% des essais publiés supposent une relation dose-toxicité monotone et croissante tandis que seulement 37.3% des papiers ne supposent qu'un ordre partiel, ce qui permet l'utilisation des designs statistiques de monothérapies. Pour sélectionner les combinaisons à retenir dans l'essai vérifiant une relation dose-toxicité monotone

croissante, les investigateurs augmentent progressivement les doses de chaque agent tout en fixant les autres. Ce processus induit un nombre limité de combinaisons à explorer et seulement un sous-ensemble de combinaisons est évalué en dépit du grand nombre de combinaisons possibles. En effet, nous avons observé que le rapport médian entre le nombre de combinaisons considérées dans l'essai et le nombre de combinaisons possibles est de 0.64 ce qui indique qu'environ un tiers de l'espace des combinaisons n'a pas été considéré. Cela signifie que les investigateurs ont choisi a priori les combinaisons à évaluer et que certaines ont été exclues. Explorer l'ensemble de l'espace de combinaisons n'est évidemment pas possible dans la pratique, et les médecins peuvent souhaiter n'explorer qu'un sous-ensemble de combinaisons. Néanmoins, le choix des combinaisons à explorer ne devrait pas être limité par la connaissance de l'ordre de toxicité, et le design devrait avoir la possibilité d'explorer n'importe quelle combinaison qu'il estime être la meilleure. En effet, en raison des interactions possibles entre les médicaments, présélectionner un sous-ensemble réduit de combinaisons induit un risque de ne sélectionner aucune combinaison dont le taux de toxicité est proche de la toxicité cible. Même si le taux de toxicités dose limitantes ciblé était de 33% dans 70% des études de notre revue bibliographique et de 16.7% dans 7.3% des études, le taux médian de toxicité associée à la dose recommandée à la fin de l'essai était beaucoup plus faible, seulement 5%. Par conséquent, les essais n'ont en général pas réussi à atteindre le taux de toxicités dose limitantes cible. Cela explique peut-être pourquoi dans 26.4% des études examinées, une combinaison intermédiaire a été rajoutée en cours d'essai induisant parfois une relation dose-toxicité non-monotone lorsqu'un agent était augmenté, tandis que l'autre était diminué. Ainsi, les méthodes pour monothérapies ne semblent pas toujours appropriées pour les essais de phase I de combinaisons lorsque plusieurs agents sont variables. Ces méthodes ne sont pas conçues pour tenir compte de la multidimensionnalité de l'espace. Plusieurs designs alternatifs, basés sur un algorithme ou sur un modèle statistique, ont été proposés par différents auteurs dans le cadre de combinaisons de deux agents. Ces méthodes donnent la possibilité d'explorer l'espace des combinaisons estimées sûres en termes de toxicité, et ont des caractéristiques opérationnelles élevées. Néanmoins, malgré le développement de nombreux modèles pour les essais cliniques de phases précoces de combinaisons, ces méthodes sont rarement utilisées dans la pratique, peut-être en raison d'un manque de compréhension de ces dernières qui nécessitent l'implication d'un expert statisticien. L'utilisation de modèles inappropriés, même dans les essais de phases précoces, peut augmenter le taux d'attrition en proposant un mauvais choix de combinaison pour les essais de phase II et de phase III.

En s'appuyant sur ce constat, la première partie de mon travail de thèse a consisté à étudier plusieurs designs de recherche de dose pour les combinaisons, représentatifs de la littérature, afin de comparer leur performance et de mettre en évidence les avantages et les inconvénients de chaque méthode. Nous avons choisi de comparer six méthodes, dont deux sont basées sur un algorithme et quatre sur un modèle. Après une large étude de simulations, il s'est avéré que

les méthodes basées sur un modèle semblent donner de meilleurs résultats que les méthodes algorithmiques lorsqu'une seule dose maximale tolérée est recommandée à la fin de l'essai. Toutes les méthodes basées sur un modèle ont des caractéristiques opérationnelles élevées avec un pourcentage de sélection correct élevé, et leurs performances sont en général comparables. Par la suite, notre objectif a donc été de proposer un design adaptatif de recherche de dose innovant pour la pratique courante qui (1) aurait de bonnes caractéristiques opérationnelles selon différents positions possibles des doses maximales tolérées, et (2) aurait en général de meilleures performances que les méthodes existantes. Nous avons décidé de modéliser la relation dose-toxicité avec un modèle logistique car ces modèles sont flexibles et souvent bien connus des médecins. Nous avons utilisé un modèle à 3 paramètres faisant intervenir chaque agent, ainsi qu'un terme d'interaction entre ces derniers. Notre processus d'allocation de doses permet d'augmenter, de diminuer ou de rester à la même combinaison, en fonction de l'estimation de probabilité de toxicité à la dose courante et de son incertitude. La recommandation de la dose maximale tolérée à la fin de l'essai est basée sur des intervalles de toxicité. En effet, la combinaison choisie pour les phases ultérieures est avec la plus grande probabilité que la (probabilité de) toxicité soit dans un intervalle autour de la toxicité cible. Selon notre étude de simulation, cette méthode semble être en mesure d'identifier la dose maximale tolérée avec un pourcentage de sélection correcte élevé sur une large variété de scénarios (avec un minimum de 56.7% et un maximum de 86.7% sur l'ensemble des scénarios). Nous avons également comparé notre méthode avec d'autres designs basés sur un modèle. Tous ceux-ci semblent être efficaces lorsque les vraies doses maximales tolérées sont situées sur la même diagonale dans l'espace des combinaisons. L'un des avantages de notre méthode par rapport aux autres est qu'elle est également performante lorsque les doses maximales tolérées ne sont pas forcément situées sur la même diagonale. Une fois de plus, pour des questions de temps, ce design a été implémenté en C/C++. Nous sommes également en train de développer un package R pour cette méthode. Nous prévoyons de terminer l'implémentation de ce package pour la soutenance de thèse.

Les agents cytotoxiques et les thérapies ciblées ont des mécanismes d'action différents, tuer les cellules cancéreuses pour les premiers et bloquer leur croissance en interférant avec les molécules spécifiques pour les seconds. Les thérapies ciblées ont émergé ces dernières années comme une autre option aux traitements cytotoxiques. Toutefois, même si certains critères permettent de valider l'action souhaitée de la molécule ciblée sur sa cible, il n'est pas immédiat que cela induise nécessairement l'activité prévue sur le cancer et entraîne donc l'efficacité attendue. Par conséquent, il peut être parfois inefficace et donc contraire à l'éthique d'administrer seul le nouvel agent ciblé. De plus, de nombreux agents cytotoxiques restent le traitement de référence pour plusieurs types de cancer. Il sera alors plutôt envisagé de le combiner et de le comparer avec le traitement cytotoxique standard. Qui plus est, le but de certaines thérapies ciblées est d'agir à l'intérieur des cellules cancéreuses afin de les inactiver. Ces cellules ne peuvent plus se répliquer, mais elles ne sont pas éliminées et restent dans le corps humain. Elles peuvent tou-

jours être observées sur différents examens pathologiques comme la radiographie ou l'imagerie par résonance magnétique, et seulement certains examens comme la tomographie par émission de positrons peuvent permettre de détecter leur inactivité. Par conséquent, l'action des agents cytotoxiques et des thérapies ciblées peut être complémentaire afin de réduire la propagation du cancer et de tuer les cellules cancéreuses restantes. Enfin, combiner plusieurs agents permet généralement de réduire les mécanismes de résistance aux médicaments. Pour toutes ces raisons, un nouveau défi dans la recherche contre le cancer est de combiner ces deux types d'agents. Pour l'instant, lorsque ces derniers sont utilisés en combinaison, les investigateurs choisissent souvent plusieurs niveaux de dose pour la nouvelle thérapie ciblée alors qu'ils fixent le niveau de dose de la chimiothérapie à la dose recommandée en monothérapie. Dans ce contexte, les designs de monothérapies pour molécule ciblée peuvent être utilisés de façon appropriée. Néanmoins, il y a un intérêt à laisser varier les deux agents car la combinaison optimale en termes de toxicité et d'efficacité combinées n'est pas nécessairement lorsque l'agent cytotoxique est à sa dose recommandée en monothérapie. Lors de la combinaison de plusieurs agents, un effet synergique sur l'efficacité est attendu. Comme l'efficacité de la molécule ciblée n'est pas monotone et croissante avec la dose contrairement à la toxicité, mais augmente puis atteint un plateau, pour la combinaison de l'agent cytotoxique et de la molécule ciblée, il n'est pas suffisant d'étudier uniquement la toxicité en tant que critère de jugement principal. Par conséquent, nous avons proposé une méthode de phase I/II permettant de combiner un agent cytotoxique avec une molécule ciblée en utilisant les caractéristiques de chaque agent. Notre objectif était de maximiser l'efficacité, tout en minimisant la toxicité sous un seuil acceptable. Nous avons supposé que la toxicité est binaire et rapidement évaluable et nous avons utilisé un modèle de régression logistique pour modéliser la probabilité de toxicité à chaque combinaison. En revanche, nous avons supposé que l'efficacité prenait plus de temps à évaluer. De façon similaire à l'analyse de survie, nous avons utilisé un modèle à hasards proportionnels avec un plateau afin de modéliser le délai d'apparition de l'efficacité. Les estimations des paramètres du modèle et les distributions des probabilités de toxicité et d'efficacité a posteriori sont mises à jour après chaque inclusion de cohortes afin de pouvoir affecter la prochaine cohorte de patients à la combinaison optimale estimée. Pour ce design, nous avons rencontré le même problème sur l'estimation du plateau et nous avons choisi d'estimer simplement le plateau à la dose ayant la probabilité a posteriori la plus élevée. En effet, en raison de la plus grande dimension de l'espace considéré, du nombre restreint de patients et de l'utilisation d'une méthode adaptative ne permettant donc pas l'exploration de tous les niveaux de dose avec suffisamment de patients, l'estimation de la localisation du plateau était difficile et sensible. Ce fut consécutivement à ce travail que nous avons décidé de développer le design de recherche de dose avec détermination du plateau dans le cadre d'une monothérapie afin de proposer une solution plus efficace à ce problème dans un contexte plus simple. Nous avons évalué notre méthode à travers une étude de simulations avec plusieurs scénarios et avons pu observer que notre méthode donne de bons résultats

en sélectionnant la combinaison optimale avec un pourcentage élevé. Cependant, les performances de notre design diminuent fortement avec le nombre de doses de la thérapie ciblée. Davantage de recherches concernant cette problématique sont nécessaires. Cette méthode a également été mis en œuvre en C/C++ et la construction d'un package R est en cours et doit être finalisée pour la fin de la thèse.

Mots-clés : Recherche de dose; Phase I-II; Oncologie; Combinaison; Cytotoxique; Molécule ciblée; Inférence bayésienne.

Abstract

Phases I are usually the first stage of testing a new drug involving human subjects. Phase I clinical trials evaluate the safety of the treatment and identify its side effects on patients with advanced cancer due to the harmfulness these treatments. The primary aim of phase I in oncology is to determine the highest dose level with an acceptable toxicity rate of the new drug on a restricted number of patients, i.e. to select a dose level with a toxicity probability closest to a given target. This recommended dose level is called maximum tolerated dose. Phase I trials are sequential dose-escalation procedures.

In recent years, unlike standard chemotherapy, targeted therapies have emerged as another type of anti-cancer agents that interact with specific molecules involved in cancer spread rather than killing cancer and healthy cells. In this context, we have developed a phase I/II dose-finding design in single-agent for molecularly targeted agents where the efficacy increases and can plateau. Our method focus on selecting the optimal dose, that is the dose associated with highest efficacy and if the plateau is reached the first dose on the plateau as it will be associated with the lowest toxicity. We used adaptive randomization in order to determine the plateau location. The proposed method gave good performance. We also extended this design on (1) unimodal relationships and (2) different biomaker's groups leading to different optimal dose in each subgroup with shared toxicity.

Methods for single agent are not appropriate for combination phase I trials as they are not designed to take into account the multi-dimensionality. We studied several existing representative methods specifically designed for combination, and compared their performance. Based on an extensive simulation study, we have noticed that model-based methods seemed to perform better than algorithm-based methods in terms of the percentage of correct combination selections when targeting a single maximum tolerated dose at the end of the trial. All model-based methods have high operational characteristics and their performances were in general comparable. On this basis, our aim was then to propose our own innovative adaptive dose-finding design that would have good operational characteristics and in general would perform better than the existing designs. We proposed a phase I dose-finding design for combination based on a logistic model with an interaction term. All the compared designs were efficient when the maximum tolerated doses were located on the same diagonal in the combination space, but the benefit of our method was that it was also efficient in other cases.

Finally a new challenge in cancer development is to combine both cytotoxic

and molecularly targeted agent. Indeed, their action can be complementary, inactivate the cells or reduce cancer growth and killing cells, but also skirt drug resistance. When combining several agents, a possible synergistic effect on the efficacy is expected. We studied both toxicity and efficacy of the combination in a Bayesian phase I/II design using the characteristics of each agent. Our goal was to maximize the efficacy while minimizing the toxicity under an acceptable threshold. We evaluated our design through a simulation study under various practical scenarios and observed that our design performed well by selecting the optimal combination with a high percentage. Nevertheless, the performance of the design highly decreased with the number of dose level of the molecularly targeted agent.

During these 3 years of PhD, we have proposed several adaptive early phase designs, either for molecularly targeted agents or combinations trials, to answer to the current need of practical statistical methods in oncology. Moreover, we have developed R packages to facilitate the access of these methods in the current practice.

Keywords: Dose-finding; Phase I-II; Oncology; Drug combinations; Cytotoxic; Molecularly Targeted Agent; Bayesian inference.