



HAL
open science

Bayesian toxicity modeling for dose regimen assessment in early phase oncology trials incorporating pharmacokinetics and pharmacodynamics

Emma Gerard

► **To cite this version:**

Emma Gerard. Bayesian toxicity modeling for dose regimen assessment in early phase oncology trials incorporating pharmacokinetics and pharmacodynamics. Statistics [math.ST]. Université Paris Cité, 2021. English. NNT : 2021UNIP5247 . tel-03872552v2

HAL Id: tel-03872552

<https://hal.science/tel-03872552v2>

Submitted on 21 May 2024

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

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

Université de Paris
**Ecole Doctorale Pierre Louis de Santé Publique : Epidémiologie et Sciences
de l'Information Biomédicale (ED393)**
Centre de Recherche des Cordeliers
Equipe HeKA : Health data and model driven Knowledge Acquisition
En partenariat avec Sanofi Research & Development, Biostatistics and Programming

Bayesian toxicity modeling for dose regimen assessment in early phase oncology trials incorporating pharmacokinetics and pharmacodynamics

Par Emma Gerard

Thèse de doctorat de Biostatistiques
Dirigée par Sarah Zohar

Présentée et soutenue publiquement le 18 novembre 2021

Devant un jury composé de :

France Mentré, Professeur, Université de Paris
Présidente de jury

Thomas M. Braun, Professor, University of Michigan
Rapporteur

Rodolphe Thiébaud, Professeur, Université de Bordeaux
Rapporteur

Sebastian Weber, Docteur, Novartis
Examineur

Christina Yap, Professor, University of London
Examinatrice

Sarah Zohar, Directrice de recherche, Université de Paris
Directrice de thèse

Marie-Karelle Rivière, Docteur, Sanofi
Co-encadrante de thèse

Moreno Ursino, Docteur, Université de Paris
Membre invité



Remerciements

Je remercie Thomas Braun et Rodolphe Thiébaud d'avoir accepté d'évaluer ce travail de thèse en tant que rapporteurs. Je remercie également France Mentré, Sebastian Weber et Christina Yap d'avoir accepté d'examiner ce travail de thèse.

Je remercie ma directrice de thèse, Sarah Zohar, pour son encadrement et son expérience dans la recherche qui ont permis de guider ce travail de thèse. Je remercie mes co-encadrants, Marie-Karelle Rivière et Moreno Ursino, pour leur encadrement, leur suivi et leur disponibilité tout au long de ma thèse. Je les remercie particulièrement pour toutes nos discussions, leurs conseils et encouragements qui m'ont permis mener cette thèse à son terme.

Je remercie Christelle Lorenzato pour l'opportunité d'avoir effectué cette thèse en partenariat avec Sanofi et pour son suivi de mon travail.

Je remercie l'ANRT et Sanofi pour le financement de cette thèse.

Je remercie Hoai-Thu Thai pour son implication dans ce travail et ses explications sur la modélisation pharmacocinétique et pharmacodynamique. Je remercie Loïc Darchy et Emmanuelle Comets pour leur retours scientifiques sur ce travail, notamment lors de mes comités de thèse.

Je remercie mes collègues au sein du département Biostatistiques et Programmation de Sanofi et au sein de l'équipe Inserm/Inria HeKA pour leur accueil.

Enfin, je remercie ma famille et mes amis pour leur soutien et leur écoute tout au long de ma thèse.

Contents

List of Figures	v
List of Tables	vii
Thesis environment	ix
Scientific production	xi
Glossary	xiii
1 Introduction	1
1.1 Drug development	1
1.2 Phase I trials in oncology	3
1.3 Dose-finding designs in phase I	6
1.4 Pharmacokinetics/dynamics in phase I	8
2 State of the art	11
2.1 Pharmacokinetic/dynamic modeling	11
2.1.1 Pharmacokinetic analysis	11
2.1.1.1 Physiological concepts: the ADME process	11
2.1.1.2 Compartmental versus non compartmental analysis	12
2.1.2 Pharmacodynamic analysis	15
2.1.3 Population pharmacokinetic/dynamic modeling	17
2.2 Dose-finding designs in oncology	20
2.2.1 Single administration	20
2.2.1.1 Standard designs evaluating a binary toxicity outcome	20
2.2.1.2 Designs evaluating an extension of the toxicity outcome	25
2.2.1.3 Designs evaluating efficacy in addition to toxicity	26
2.2.1.4 Designs evaluating drug combinations	27
2.2.2 Multiple administrations	28
2.2.2.1 Dose-schedule finding as a two-dimensional problem	28
2.2.2.2 Modeling the time to DLT after multiple administrations	30
2.2.2.3 Evaluating the sequence of administrations	32
2.3 Dose-finding designs incorporating pharmacokinetics	34
2.3.1 Single administration	34
2.3.2 Multiple administrations	35

3	Objectives of the thesis	39
3.1	Motivating trial	39
3.2	Objectives and outline of the manuscript	44
4	Bayesian dose regimen assessment using PK/PD	45
4.1	Summary	46
4.2	Publication	48
5	Bayesian modeling of a bivariate toxicity for dose regimens	63
5.1	Summary	64
5.2	Publication	66
6	Conclusion and discussion	87
	Bibliography	93
	Abstract	107
	Résumé	109
	Résumé long	111
A	Bayesian inference	119
A.1	Bayesian paradigm	119
A.2	Computation of the posterior distribution	121
B	PK/PD models and estimation	123
B.1	Details on the PK/PD models	123
B.2	Estimation of population parameters for nonlinear mixed effects models	128
C	Supporting information: Bayesian dose regimen assessment using PK/PD	131
D	Supporting information: Bayesian modeling of a bivariate toxicity for dose regimens	157

List of Figures

1.1	Drug development divided in five successive steps.	1
1.2	Three examples of dose regimens during three cycles of treatment. . .	3
1.3	Example of dose-toxicity and dose-efficacy curves.	5
1.4	Probability of success in oncology compared to other therapeutic ar- eas from 2011 to 2020.	7
1.5	Pharmacokinetic and pharmacodynamic modeling.	8
2.1	Illustration of PK.	12
2.2	One and two compartment(s) models.	13
2.3	Drug concentration after a single administration.	14
2.4	Drug concentration after multiple administrations.	15
2.5	Indirect response models.	16
2.6	PD response after a single administration.	17
2.7	Population PK modeling.	19
2.8	3+3 design.	20
2.9	Continual reassessment method (CRM).	22
2.10	Decision rules for a model-assisted design.	24
2.11	Partial order in the combination of two agents.	27
2.12	Partial order in dose-schedule finding as a two-dimension problem. . .	29
2.13	Triangular and Weibull hazard functions	31
3.1	Schematic representation of SAR440234.	39
3.2	Cytokine model.	42
3.3	Example of a panel of dose regimens defined with intra-patient dose- escalation.	43
3.4	Modeling proposal.	44
4.1	Trial scheme: the DRtox method is applied at the end of the dose- escalation stage of a phase I trial	51
4.2	Concentration (up) and cytokine (down) profiles of two patients . . .	55
4.3	The first three subplots represent the panel of dose regimens. In the last subplot in the lower right corner, the dose regimen toxicity relationship is represented for each scenario	56
4.4	Violin plots of the estimated toxicity probabilities in an additional sce- nario in which the dose regimen panel missed the true MTD-regimen and in Scenario 3	58
5.1	Illustration of the proposed modeling process of the bivariate toxicity endpoint at the end of the dose-escalation phase of the trial	69

5.2	Definitions of the six toxicity scenarios in terms of the probabilities of dose limiting toxicity (DLT), cytokine release syndrome (CRS), and other DLT (DLT_o).	75
5.3	Violin plots of the estimated probabilities of dose limiting toxicity (DLT) in Scenario 1.	79
5.4	Inclusion process of the simulated trial under a modified continual reassessment method (CRM).	79
6.1	Exemple de courbes de toxicité et d'efficacité	112
6.2	Proposition de modélisation	113
6.3	Effet de l'escalade de dose intra-patient sur la réponse PK/PD	114
A.1	Illustration of Bayes' rule	120
B.1	Cytokine model.	123
B.2	Concentration profiles.	126
B.3	Cytokine profiles using the initial parameters.	127
B.4	Cytokine profiles using the modified parameters.	127

List of Tables

1.1	Toxicity grades according to CTCAE categorization for two specific adverse events.	4
2.1	Toxicities and severity weights.	25
3.1	Revised CRS grading.	40
3.2	Example of a panel of dose regimens with intra-patient dose-escalation.	40
3.3	Example of a 3+3 design applied on the panel of dose regimens defined with intra-patient dose-escalation.	41
4.1	Definition and values of the PK/PD parameters used for the simulation study	54
4.2	Proportions that each dose regimen is being selected as the MTD-regimen	57
5.1	Dose regimens defined in Set A and Set B used in the simulation study	73
5.2	Proportions of selecting each dose regimen \mathbf{S}_k as the MTD-regimen over the 1000 trials in the six main toxicity scenarios	76
5.3	Proportions that each dose regimen is being selected as the MTD-regimen	77
B.1	Estimated PK/PD parameters for blinatumomab and PK/PD parameters used for the simulations.	125
B.2	Example of a panel of dose regimens defined with intra-patient dose-escalation.	126

Thesis environment

In France, the Cifre (*Conventions Industrielles de Formation par la REcherche*) system, that stands for industrial agreement of training through research, aims to promote the development of public-private partnership by the supervision of a PhD student. In a Cifre system, managed by the ANRT (*Association Nationale Recherche Technologie*) that stands for national agency for research and technology, a company hires a PhD student to conduct his/her research project in partnership with a French laboratory and the company receives in return a grant by the ANRT.

In my case, this PhD was conducted with the Cifre grant number 2018/0530 in partnership with Sanofi-Aventis R&D and the HeKA (Health data and model driven knowledge acquisition) research team (Inserm, Inria). In Sanofi, I was working in the Biostatistics & Programming department, both in the methodological and oncology teams, and supervised by Marie-Karelle Rivière. In the HeKA team, located in the *Centre de recherche des Cordeliers*, I was directed by Sarah Zohar and co-supervised by Moreno Ursino.

Scientific production

Published papers:

- Gerard E, Zohar S, Lorenzato C, Ursino M, Riviere M-K. Bayesian modeling of a bivariate toxicity outcome for early phase oncology trials evaluating dose regimens. *Statistics in Medicine* 40, 23 (2021), 5096-5114.
<https://doi.org/10.1002/sim.9113>
- Gerard E, Zohar S, Thai H-T, Lorenzato C, Riviere M-K, Ursino M. Bayesian dose regimen assessment in early phase oncology incorporating pharmacokinetics and pharmacodynamics. *Biometrics*. 2021. Online, ahead of print.
<https://doi.org/10.1111/biom.13433>

Oral presentations at conferences:

- ISCB (International Society for Clinical Biostatistics), Lyon 2021: Bayesian modeling of a bivariate toxicity outcome for early phase trials evaluating dose regimens
- Epiclin, Marseille 2021: Modélisation Bayésienne d'une toxicité bivariée en essais de phase précoce en oncologie évaluant des régimes de doses
- ISCB (International Society for Clinical Biostatistics), Krakow 2020: Dose-sequence finding using pharmacokinetics and pharmacodynamics
- IBC (International Biometric Conference), Seoul 2020: Dose-sequence finding using pharmacokinetics and pharmacodynamics

Glossary

- ADME: Absorption, Distribution, Metabolism, Excretion
- AUC: Area under the curve
- BiTE: Bispecific T cell engager
- BOIN: Bayesian optimal interval design
- CRM: Continual reassessment method
- CRS: Cytokine release syndrome
- CTCAE: Common terminology criteria for adverse events
- DLT: Dose limiting toxicity
- DLT_o: Other DLT (different from CRS)
- DRtox: Dose regimen assessment method
- EMA: European medicines agency
- ESS: Effective sample size
- EWOC: Escalation with overdose control
- FDA: Food and drug administration
- IFN: Interferons
- IL: Interleukins
- MCMC: Markov chain Monte Carlo
- MTD: Maximum tolerated dose
- MTD-regimen: Maximum tolerated dose regimen
- NCI: National cancer institute
- PCS: Proportion of correct selection
- PK/PD: Pharmacokinetics/Pharmacodynamics
- POCRM: Partial order continual reassessment method
- RP2D: Recommended phase II dose

- SAEM: Stochastic approximation expectation maximization
- TITE-CRM: Time-to-event continual reassessment method
- TPI: Toxicity probability interval

Chapter 1

Introduction

1.1 Drug development

Drug development is a very long and complex research process carried out under strong surveillance of health authorities, for example the Food and Drug Administration (FDA) in the United States of America and the European Medicines Agency (EMA) in the European Union. The aim of drug development is to validate that the new drug is safe and effective before entering the marketing phase. According to the FDA [4], drug development follows five successive steps as illustrated in Figure 1.1.




Figure 1.1 – Drug development divided in five successive steps.

Step 1: Discovery Thousands of molecules are screened to detect if there exist any promising ones that are worth being investigated in further steps of the research process. First information on the mechanism of action, dosage, mode of administration or interaction with other drugs is gathered.

Step 2: Preclinical research The new drug is assessed for activity, toxicity and pharmacology during in-vitro (tube, cells culture), in-vivo studies (in animals such as rodents, non-human primates etc.) or in silico using computer simulations. Depending on the results obtained, the new drug can continue the development

process in humans. Data collected during this phase can help design the subsequent studies.

Step 3: Clinical research During this step, lasting about 7 to 10 years, the new drug is administered to humans during clinical trials that are usually divided in 3 phases.

Phase I trials are the first stage of human experimentation and aim at evaluating the safety of the drug in a small number of subjects (less than 100). They are either carried out on healthy volunteers (for most therapeutic areas) or on patients due to the harmfulness of some drugs (in oncology for example).

Phase II trials evaluate the efficacy of the drug on one or multiple doses. They are carried out on up to several hundreds of patients with the targeted disease. Phase I/II trials are also conducted to combine the safety and efficacy assessment in a single study.

Finally, phase III trials (or confirmatory trials) are performed on hundreds to several thousands of patients to demonstrate the efficacy of the new drug compared to the standard of care according to prespecified hypotheses. Seamless phase II/III trials can also be performed in order to combine the phase II and phase III in the same trial [143].

Step 4: Review of health authorities After the phase III trial, if the new drug shows evidence of tolerance and efficacy for the disease of interest, the pharmaceutical company can file an application to market the drug. Then, the health authorities examine all the results obtained to decide if the drug should be approved. If the drug receives the market authorization, it can then be distributed to the public.

Step 5: Post market safety monitoring After obtaining the approval, phase IV studies are carried out on thousands of patients to study the drug in “real life” conditions. One of the objectives of these studies is to detect rare or long term toxicities that have not been observed during the clinical trials.

The work of this thesis focused on phase I trials in oncology that will be further developed in the following section.

1.2 Phase I trials in oncology

Phase I trials in oncology are the first stage in human experimentation of the drug. They are performed on a small number of patients who have usually been heavily treated and for whom standard therapies have failed. These trials aim to evaluate the toxicity profile of a new drug or of a combination of drugs in order to define a tolerable dose [118]. They also aim at describing the pharmacokinetic and pharmacodynamic profiles of the drug (details will be given in section 1.5). During phase I/II trials, the objective is also to demonstrate the anti-tumor activity, namely the efficacy, of the drug.

Anti-cancer therapies are usually given in multiple treatment cycles, where a cycle is defined by the National Cancer Institute (NCI) as “a period of treatment followed by a period of rest (no treatment) that is repeated in a regular schedule” [12]. The first cycle begins at the first administration of the drug, and the duration of the cycle is specified in the trial protocol. For example, in Figure 1.2, a cycle is defined as three weeks of therapy.

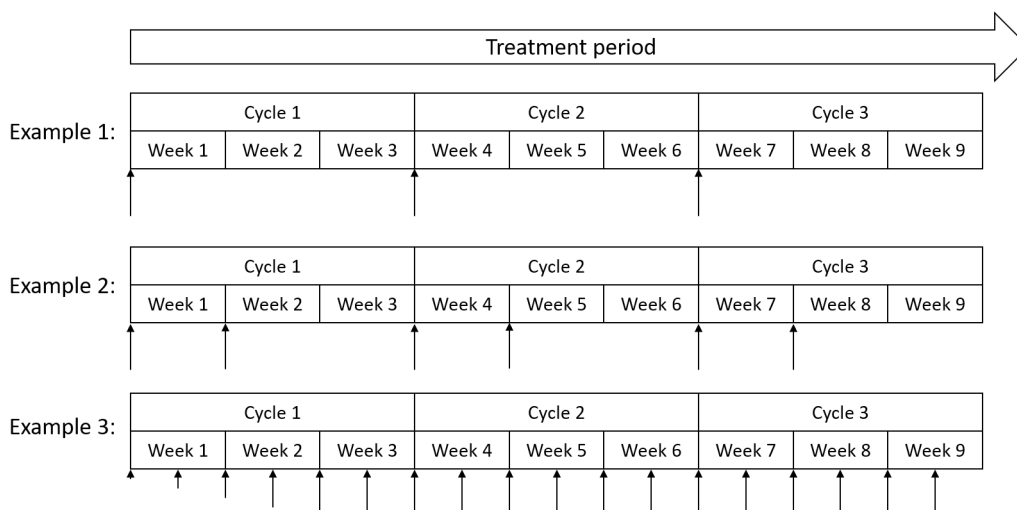


Figure 1.2 – Three examples of dose regimens during three cycles of treatment, where a cycle corresponds to three weeks of therapy. The arrows represent drug administrations, their heights are proportional to the amount of administered drug.

During a cycle, the drug is given once or repeatedly according to the treatment schedule defined by the NCI as “a step-by-step plan of the treatment that a patient is going to receive. A treatment schedule includes the type of treatment that will be given (such as chemotherapy or radiation therapy), how it will be given (such as by mouth or by infusion into a vein), and how often it will be given (such as once a day or once a week). It also includes the amount of time between courses of treatment and the total length of time of treatment” [13]. For example, in Figure 1.2, the drug is given once per cycle in example 1, twice per cycle in example 2 (on days 1 and 8 of each cycle) and six times per cycle in example 3 (twice per week).

In a cycle, the same dose can be given repeatedly but the dose can also vary between the administrations, according to the trial protocol. In Figure 1.2, the same dose is administered in examples 1 and 2. However, in example 3, the dose is slowly increased during the first two weeks of cycle 1, and the highest dose achieved is then

repeated from the last week of cycle 1 until the end of therapy. Treatment regimen is defined by the NCI as “a treatment plan that specifies the dosage, the schedule, and the duration of treatment” [11]. Throughout this thesis, we will consider the term dose regimen that we define as the combination of the schedule of administration and the dose, that can vary during the administrations of the drug. Moreover, the protocol of the trial usually also describes the plan to modify the dose given in subsequent cycles according to the patient’s history, in particular in case of adverse events. For example, the dose can be reduced in subsequent cycles if the patient experiences too severe toxicity.

Toxicity is measured by the occurrence of adverse events defined by the NCI as “an unexpected medical problem that happens during treatment with a drug or other therapy” [9]. The NCI Common Terminology Criteria for Adverse Events (CTCAE) [8] provides a standardized classification of adverse events related to anti-cancer therapy according to their severity in five severity grades, from grade 1, which is a mild adverse event, to grade 5, which is death. Examples of this classification for two specific adverse events, hepatic hemorrhage and cytokine release syndrome, are provided in Table 1.1.

	Hepatic hemorrhage: A disorder characterized by bleeding from the liver	Cytokine release syndrome: A disorder characterized by fever, tachypnea, headache, tachycardia, hypotension, rash, and/or hypoxia caused by the release of cytokines
Grade 1: Mild	Mild symptoms; intervention not indicated	Fever with or without constitutional symptoms
Grade 2: Moderate	Moderate symptoms; intervention indicated	Hypotension responding to fluids; hypoxia responding to <40% O ₂
Grade 3: Severe	Transfusion indicated; invasive intervention indicated; hospitalization	Hypotension managed with one pressor; hypoxia requiring \geq 40% O ₂
Grade 4: Life-Threatening	Life-threatening consequences; urgent intervention indicated	Life-threatening consequences; urgent intervention indicated
Grade 5: Death	Death	Death

Table 1.1 – Toxicity grades according to CTCAE categorization for two specific adverse events.

For each patient, the type, severity grade and duration of each adverse event occurring during the therapy are collected. In phase I trials, all adverse events occurring during a prespecified observational window (usually during the first cycle) are summarized in a single binary outcome, defined as the dose-limiting toxicity

(DLT), which denotes the presence or absence of severe adverse events (that should be avoided). The exact definition of DLT varies according to the type of therapy and is explicitly provided in the protocol of the trial. For example, a DLT can be defined as the occurrence of an adverse event of grade 3 or more as well as an adverse event of grade 2 having a long duration during the first cycle of therapy.

The main objective of phase I trials is to determine the maximum tolerated dose (MTD) which can be defined as the dose having an estimated probability of DLT (during cycle 1) the closest to a predefined target toxicity level (for example, 0.30) [118]. The MTD is usually defined from an initial panel of dose-levels. The rationale for finding the MTD is linked to the assumption, on cytotoxic agents, that the probabilities of toxicity and efficacy monotonically increase with the dose, as illustrated in Figure 1.3 [84]. Therefore, a dose with a DLT probability close to the target should exhibit a higher chance of anti-tumor activity compared to a dose having a very low risk of DLT and toxicity may be viewed as a surrogate endpoint for efficacy [40]. A trade-off between toxicity and efficacy is then researched as the MTD exhibits the highest probability of efficacy among the doses with an acceptable probability of toxicity. For example, in Figure 1.3, dose-level 4 is the estimated MTD, in the panel of six dose-levels, having a probability of toxicity of 0.30 and a probability of efficacy of almost 0.60. The recommended dose for further evaluation, the so called recommended phase II dose (RP2D), is then defined as either the MTD or one of the doses below the MTD [83].




Figure 1.3 – Example of dose-toxicity and dose-efficacy curves.

Specific dose-finding designs have been developed for phase I trials in order to determine the MTD under ethical concerns.

1.3 Dose-finding designs in phase I

Phase I designs in oncology have been developed to find the MTD under ethical concerns [84, 42]. Firstly, they should include a small number of patients (who usually are patients for whom standard therapies have failed). Secondly, they should avoid to expose patients to high toxicity and should therefore minimize the number of patients receiving doses that are too toxic. Finally, phase I trials in oncology also have a therapeutic aim for the patients included, and therefore dose-finding designs should also minimize the number of patients receiving ineffective doses.

In the example shown in Figure 1.3, dose-levels 1 and 2 are safe but ineffective (probability of efficacy lower than 0.20) while dose-levels 5 and 6 exhibit a high probability of efficacy but are too toxic (probability of toxicity higher than 0.40). In this example, the dose-finding design should maximize the number of patients treated around dose-level 4. In this context, randomizing the patients among the panel of doses is not acceptable as many patients would receive either ineffective or too toxic doses.

To deal with these ethical requirements, sequential dose-escalation procedures have been proposed where patients are sequentially enrolled in the trial by cohorts of small sample size (for example of size 3). The first cohort of patients is usually treated at the lowest dose, and successive cohorts are progressively treated at higher doses until the MTD is found. The doses given in the successive cohorts are determined with either an algorithm based approach or a model based approach from the data observed on all previous cohorts. The design should be defined so that dose-escalation does not occur too quickly, to avoid too toxic doses, and not too slowly, to avoid ineffective doses. Standard dose-finding designs will be detailed in Chapter 2.

According to the 2021 report on clinical development success rates between 2011 and 2020, oncology had the third lowest probability of success in comparison with 14 other indications. The probability of success was 5.3% in oncology versus 9.3% in the other indications, where the highest probability of success was in hematology trials (23.9%) [18]. To further understand this low probability of success, the report compared the probability of success of every clinical transition of oncology trials versus non-oncology trials, these probabilities are represented in Figure 1.4. Oncology had a lower success rate in comparison to the 14 non-oncology therapeutic areas in almost every clinical transition, except for submission to approval with 92% compared to 90.2%. Oncology had the lowest probability of success in phase III with a value of 47.7% while this probability was equal to 61.3% for other indications.

This high failure rate in oncology was not necessarily due to the development of ineffective drugs, but could also be related to a poor identification of the right dose (or dose regimen) in early phase trials. Therefore, the choice of the RP2D in phase I is crucial for success in further stages of the trials and the development of innovative dose-finding designs are a key to reliably estimate the most promising dose (or dose regimen) to be studied in phase II and III.

Usual dose-finding designs have been developed to determine the MTD (and the RP2D) after one cycle of therapy, under a specific schedule and route of adminis-

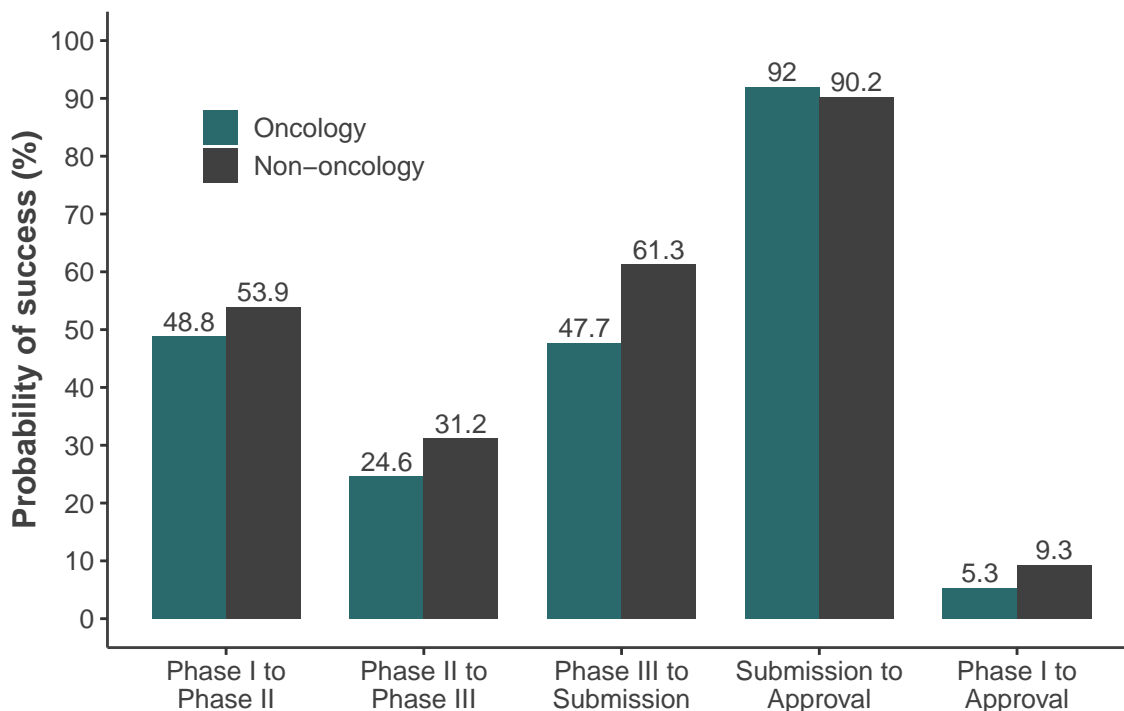


Figure 1.4 – Probability of success in oncology compared to other therapeutic areas from 2011 to 2020 [18].

tration, for example an intravenous infusion every two weeks. However, varying the schedule of administration in addition to the dose (defined as the per-administration dose or the cumulative dose over the cycle) or considering alternative dose regimens (with intra-patient dose-escalation for example) could improve treatment safety, while preserving future efficacy. In this context, the use of pharmacokinetic and pharmacodynamic information in addition to standard dose-finding methods could improve the choice of the RP2D and the success rate of the drug in further stages of the development.

1.4 Pharmacokinetics/dynamics in phase I

Clinical pharmacology is based on the fact that the intensity of many pharmacological effects is linked to the amount of drug in the body [53]. Pharmacokinetics (PK) studies the relationship between the dose administered and the concentration of drug over time, while pharmacodynamics (PD) studies the relationship between the concentration of drug and the biochemical or physiologic effect. To simplify, PK studies “what the body does to the drug” while PD studies “what the drug does to the body”. Pharmacokinetic/pharmacodynamic (PK/PD) modeling therefore describes the relationship between the dose regimen and the effects observed as illustrated in Figure 1.5.




Figure 1.5 – Pharmacokinetic and pharmacodynamic modeling, inspired by [53].

Drug concentration is measured from regular blood samples after drug administration, and PK aims at studying the evolution of drug concentration over time. The PD marker is also assessed regularly when PD analyses are performed. The timings of PK/PD criteria assessments are prespecified in the trial protocol.

PK/PD modeling can be performed at different stages of the clinical development to evaluate the dose-exposure-response relationship. According to the FDA guidance for industry, PK modeling can be performed to predict the drug exposure associated with doses or dosing regimens that have not yet been investigated in previous clinical trials [156]. In particular, “a population PK analysis can be used to predict PK changes resulting from the inclusion of a loading dose, changing the dose, or altering the dosing frequency of a dosing regimen for later trials in the drug development program.”

PK/PD modeling can firstly be performed in preclinical trials which provide first information on toxicity and activity in various species. This first modeling can support the choice of the initial safe dose to be tested in phase I trials, the dose range, the schedule of administration and the optimal outcomes samplings required for PK/PD estimation [61]. Traditionally, the starting dose was defined from preclinical trials as one tenth of the lethal dose in 10% of rodents, or one sixth of the highest non-severely toxic dose in nonrodents (for example in non-human primates) [71]. However, for high-risk medicinal products, the EMA recommends to determine the minimal anticipated biological effect level from PK/PD modeling using all in vitro and in vivo information and considering the inter species differences [49, 28].

In phase I trials, characterizing the PK profile of the drug is usually defined as a secondary objective while the PD profile is an exploratory objective. PK modeling can support the choice of the recommended dose or dose regimen for phase II trials based on a targeted exposure to the drug and can also evaluate the effect of varying the treatment schedule. However, in a survey performed by Comets and Zohar [46], PK results were often described separately from the dose-finding results and few attempts were made to link them together.

In phase I trials, toxicity information could therefore be combined to PK/PD information, into a unified modeling approach, to support the choice of the dose or dose regimen for further stages of the drug development.

Chapter 2

State of the art

2.1 Pharmacokinetic/dynamic modeling

As mentioned in the introduction, pharmacokinetics/pharmacodynamics (PK/PD) refers to the relationship between the dose regimen administered and the effects observed.

2.1.1 Pharmacokinetic analysis

PK studies the relationship between the dose administered and the concentration of drug in the body over time, and relies on physiological concepts.

2.1.1.1 Physiological concepts: the ADME process

PK is based on the course of the drug in the body over time. Indeed, the drug needs to cross restrictive barriers defined by the absorption, distribution, metabolism and excretion process (ADME) [36].

- **Absorption:** The absorption denotes the transfer of the drug from the site of administration into the bloodstream. Depending on the route of administration (intravenous, intramuscular injection, oral ingestion, etc.), the drug needs multiple steps before reaching the bloodstream. Therefore, only a fraction of the initial dose administered successfully arrives in the bloodstream, which is defined by the bioavailability factor. In case of an intravenous injection, the bioavailability factor is 1 as the dose is directly injected into the bloodstream. However, if the drug is taken orally, the bioavailability is reduced mainly due to the digestive system.
- **Distribution:** The distribution denotes the transfer of the drug from the bloodstream to the effective sites. The volume of distribution (V) is a fictive volume in which the drug would be distributed to obtain the observed concentrations (this value can be higher than the total volume of the organism).
- **Metabolism:** The metabolism denotes the chemical reactions that convert the drug into compounds that are more easily eliminated. For first-order kinetics, the amount of drug metabolized per unit of time is proportional to drug concentration.

- **Excretion:** The excretion denotes how the drug is eliminated from the body essentially through the kidneys. The excretion is defined by the clearance parameter (Cl), analogous to the creatinine clearance [128], defined as the volume of plasma that is completely cleared of the drug in a unit of time. The half-life ($t_{1/2}$) denotes the time interval when drug concentration is reduced to half of its value.

PK aims to analyze the evolution of drug concentration over time. Drug concentrations are obtained from regular samplings of the blood. The PK of a drug can be described by some measures of exposure such as the area under the curve (AUC), the maximum concentration achieved (C_{\max}), the time to reach the maximum concentration (t_{\max}) or the half-life of the drug ($t_{1/2}$) as defined above. These PK values are represented in the left part of Figure 2.1.




Figure 2.1 – Illustration of PK where the left part of the figure represents PK concepts and the right part illustrates PK samplings. In the right part, the points illustrate the samplings, the plain line represents the concentration curve obtained from simply linking the samplings and the dashed line represents the model fitted.

2.1.1.2 Compartmental versus non compartmental analysis

The non-compartmental analysis consists in estimating directly the dose concentration curve from the observed concentrations obtained from the blood samples, as illustrated in the right part of Figure 2.1. Most PK parameters can then directly be obtained (C_{\max} , t_{\max}) and the AUC can be estimated using the trapezoidal rule. This analysis requires no assumption on the ADME process but needs a lot of concentration measures to have a precise dose concentration curve over time. However,

this approach cannot study the PK of the drug when varying the amount of the drug administered or the schedule of administration.

The compartmental analysis describes the relationship between the drug concentration and time with an explicit mathematical model through assumptions on the ADME process. The organism is simplified as a system of theoretical compartments, where a compartment is defined as a kinetically homogeneous region of the body [110]. The simplest model only contains one compartment, the central compartment, while more complex models include one or two peripheral compartment(s) linked to the central one. One and two compartment(s) models with a linear elimination process are illustrated in Figure 2.2. The model should be precise enough to describe the evolution of drug concentration over time but should remain simple enough to have a reasonable number of parameters to estimate (that may have a biological interpretation) and to avoid overfitting.




Figure 2.2 – One (a) and two (b) compartment(s) models. V (V_2) is the volume of distribution of the central (second) compartment, k is the elimination rate constant, and k_{12} (k_{21}) is the distribution rate constant from compartment 1 (2) to compartment 2 (1).

With a compartmental model, drug concentration can be mathematically defined. For a one compartment model with linear elimination, let the clearance of elimination be defined as $Cl = kV$, where k is the elimination rate constant and V is the volume of distribution. After a single administration of dose D at time t_D , drug concentration depends on the route of administration as provided below.

For an intravenous bolus administration, drug concentration is defined as follows:

$$C(t) = \frac{D}{V} \exp(-k(t - t_D)) \quad (2.1)$$

For an intravenous infusion of duration T_{inf} , drug concentration is defined as follows:

$$C(t) = \begin{cases} \frac{D}{T_{inf}} \frac{1}{kV} (1 - \exp(-k(t - t_D))) & \text{if } t \leq T_{inf} \\ \frac{D}{T_{inf}} \frac{1}{kV} (1 - \exp(-kT_{inf})) \exp(-k(t - t_D - T_{inf})) & \text{if } t > T_{inf} \end{cases} \quad (2.2)$$

For an oral administration, let k_a be the absorption rate constant and drug concentration is defined as follows:

$$C(t) = \frac{D}{V} \frac{k_a}{k_a - k} (\exp(-k(t - t_D)) - \exp(-k_a(t - t_D))) \quad (2.3)$$

The effect of the route of administration on the drug concentration is illustrated in the left part of Figure 2.3.




Figure 2.3 – Drug concentration after a single administration for various PK models. In the left part of the figure, the impact of the route of administration is illustrated, while the difference between a 1-compartment and a 2-compartment model is shown in the right part. The figures were generated with $k = 0.05 \text{ h}^{-1}$, $V = 10 \text{ L}$, $D = 100 \text{ mg}$, $T_{\text{inf}} = 10 \text{ h}$, $k_a = 0.25 \text{ h}^{-1}$, $k_{12} = k_{21} = 0.1 \text{ h}^{-1}$.

These models can also include an effect compartment to delay the response to the drug.

After n bolus of doses $\{D_1, \dots, D_n\}$ administered at times $\{t_{D_1}, \dots, t_{D_n}\}$, drug concentration is defined as follows:

$$C(t) = \sum_{i=1}^n \mathbf{1}_{\{t > t_{D_i}\}} \frac{D_i}{V} \exp(-k(t - t_{D_i})) \quad (2.4)$$

Similar equations for the other routes of administration can be developed [25].

The assumption of linear elimination implies that the AUC of the drug is not affected by different schedules of administration, if the same total dose is administered [128]. For example, the total AUC after a single administration of 100 mg is the same as after two administrations of 50 mg administered 24 or 72 hours apart, as illustrated in Figure 2.4.

Mathematical expressions for the concentration with additional compartments or under a different assumption for the elimination can also be obtained [25].

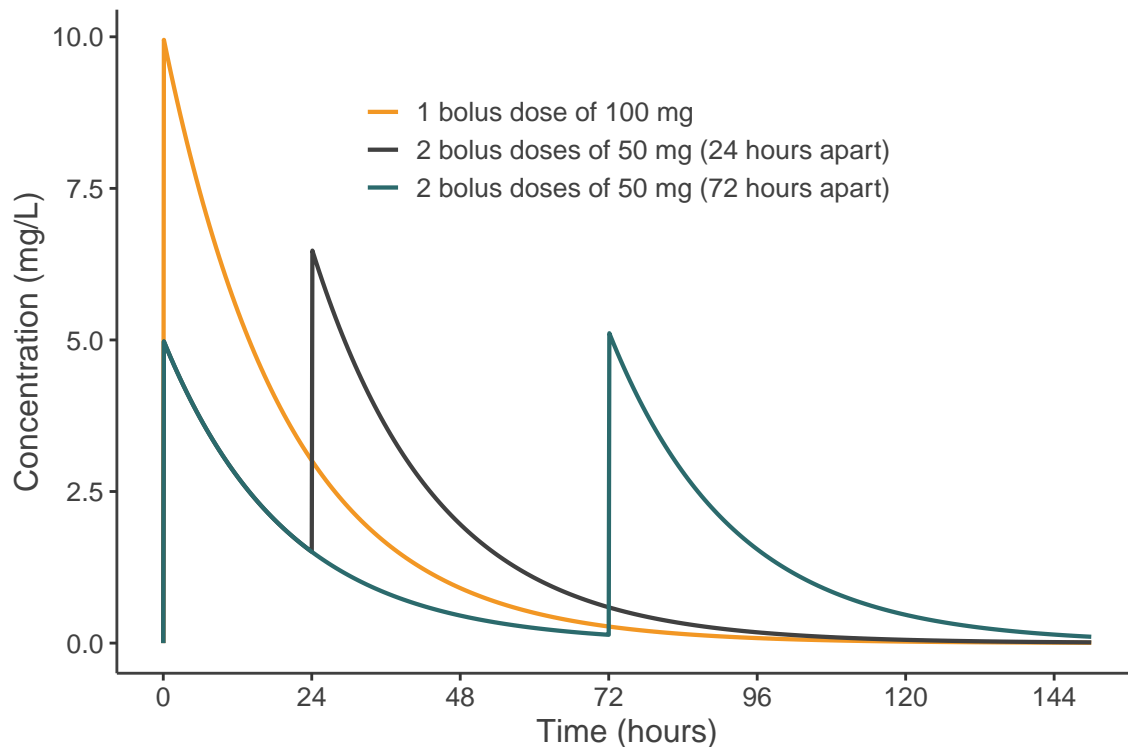


Figure 2.4 – Drug concentration after multiple administrations using a one compartment model with linear elimination. The total AUC is the same for the three dose regimens. The figure was generated using $k = 0.05 \text{ h}^{-1}$ and $V = 10 \text{ L}$.

2.1.2 Pharmacodynamic analysis

PD aims to evaluate the relationship between drug concentration and the biochemical or physiologic effect, defined either as a continuous or discrete outcome. We will focus on a continuous PD outcome, for example the response of a PD biomarker related to toxicity or efficacy.

Direct response models assume that drug concentration is directly linked with the PD response. The sigmoid Emax model can be expressed as follows:

$$E(t) = E_0 + E_{\max} \frac{C(t)^\gamma}{C(t)^\gamma + EC_{50}^\gamma}, \quad (2.5)$$

where E_0 is the baseline effect prior to drug administration, E_{\max} is the maximum effect, EC_{50} is the drug concentration to reach half of the maximum effect and γ is the Hill coefficient that allows flexibility in the sigmoid curve. The simple Emax model is obtained assuming $E_0 = 0$ and $\gamma = 1$. Extensions of this model include the inhibition model or the addition of an effect compartment to delay the PD response.

Indirect response models assume that drug concentration does not directly affect the PD response but impacts either the production (k_{in} parameter) or the elimination (k_{out} parameter) of the PD response by stimulation or inhibition, as represented in Figure 2.5 [135].

The four basic indirect response models are defined as follows:




Figure 2.5 – Indirect response models where k_{in} (k_{out}) is the constant for production (elimination) of the effect. Drug concentration can stimulate (+) or inhibit (–) the production or elimination of the response.

- Stimulation of k_{in}

$$\frac{dE(t)}{dt} = k_{in} \left(1 + \frac{S_{\max}C(t)}{SC_{50} + C(t)} \right) - k_{out}E(t) \quad (2.6)$$

- Inhibition of k_{in}

$$\frac{dE(t)}{dt} = k_{in} \left(1 - \frac{I_{\max}C(t)}{IC_{50} + C(t)} \right) - k_{out}E(t) \quad (2.7)$$

- Stimulation of k_{out}

$$\frac{dE(t)}{dt} = k_{in} - k_{out} \left(1 + \frac{S_{\max}C(t)}{SC_{50} + C(t)} \right) E(t) \quad (2.8)$$

- Inhibition of k_{out}

$$\frac{dE(t)}{dt} = k_{in} - k_{out} \left(1 - \frac{I_{\max}C(t)}{IC_{50} + C(t)} \right) E(t) \quad (2.9)$$

$S_{\max} > 0$ is the maximum stimulation factor to affect k_{in} or k_{out} , SC_{50} is the drug concentration that produces half of the maximum stimulation, $0 < I_{\max} \leq 1$ is the maximum inhibition factor to affect k_{in} or k_{out} and IC_{50} is the drug concentration that produces half of the maximum inhibition.

Similarly to the direct response model, these indirect response models have been extended, for example to model a tolerance or rebound phenomena by defining a precursor [134]. Some PD responses are illustrated in Figure 2.6.

Other mechanistic models based on ordinary differential equations can be developed to describe various biomarkers or tumor size kinetics for example [122]. Although mechanistic PK/PD models are based on biological knowledge, they attempt to remain relatively parsimonious and might miss some signals [29]. Quantitative and systems pharmacology can be seen as an extension of PK/PD modeling as it integrates knowledge from various disciplines (classic pharmacology, chemical biology, structural biology, applied mathematics, medicine) in order to understand precisely how drugs modulate cellular networks in space and time and how they impact human pathophysiology [141]. In this manuscript, we will focus on PK/PD models.

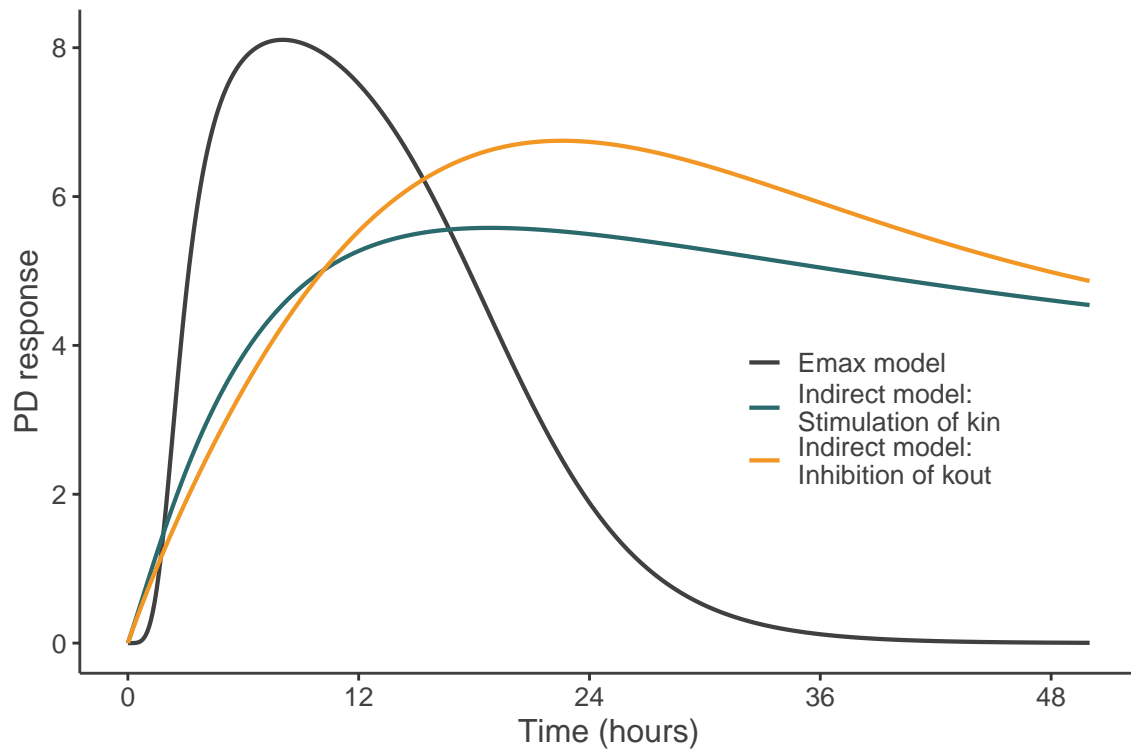


Figure 2.6 – PD response after a single administration. The figure was generated with $E_0 = 0$, $E_{\max} = 10$, $EC_{50} = 5$ mg/L, $\gamma = 5$, $k_{in} = 0.75$ h⁻¹, $k_{out} = 0.2$ h⁻¹, $S_{\max} = 1$, $SC_{50} = 5$ mg/L, $IC_{50} = 5$ mg/L, $I_{\max} = 1$. Drug concentration was obtained with a 1-compartment oral model where $k = 0.05$ h⁻¹, $V = 10$ L, $D = 100$ mg, $k_a = 0.25$ h⁻¹.

2.1.3 Population pharmaco-kinetic/dynamic modeling

Multiple observations of drug concentration and PD response are obtained for each patient over time (longitudinal data). Sampling times are specified in the trial protocol but may vary among individuals due to patients' follow-up. PK and PD data are characterized by an important inter-patient variability due to intrinsic (genetic polymorphisms, renal or liver impairment) and extrinsic (food consumption, concomitant medication) factors, and measurement errors [156].

Individual PK/PD estimation can be performed on each patient data. It requires a lot a data for each patient and does not attempt to share information between patients. Non compartmental analysis can be performed to estimate PK/PD parameters without assumption on the underlying model, or models defined above can be fitted using non linear regression models. In the right part of Figure 2.1, the plain line represents non-compartmental analysis and the dashed line represents individual fitting.

Population PK/PD modeling was proposed to aggregate all individual data in one unified model. The objective is to estimate the population PK/PD profile (defined as the profile of a standard patient) and the individual PK/PD profiles accounting for inter and intra variability. The population modeling approach, recommended by both the FDA and the EMA in specific guidances [156, 47], requires

less individual data and can be used for predictions and simulations, for example predictions of drug exposures following alternative dose regimens. However, population approaches are more computationally intensive due to the use of nonlinear mixed effect models.

Let N be the number of patients in the trial. Let $\mathbf{y}_i = (y_{i1}, \dots, y_{in_i})$ be the measurements of patient $i \in \{1, \dots, n\}$ having n_i observations. Observation y_{ij} measured at time t_{ij} is assumed to follow the model defined below:

$$y_{i,j} = f(\boldsymbol{\theta}_i, \mathbf{x}_{i,j}) + g(\boldsymbol{\theta}_i, \mathbf{x}_{i,j}, \boldsymbol{\xi}) \epsilon_{ij}, \epsilon_{ij} \sim \mathcal{N}(0, 1). \quad (2.10)$$

The structural model, f , is a non linear function that represents the biological process generating the PK/PD data. f can be explicitly defined (1 compartment PK model for example) or be the solution of a differential equation (indirect response model for example). $\boldsymbol{\theta}_i$ represents the individual parameters of patient i and $\mathbf{x}_{i,j}$ represents the descriptive variables (sampling time and the dose for the 1 compartment bolus model, concentration for the Emax model, etc.).

For example, for a single bolus administration, $f(\boldsymbol{\theta}_i, \mathbf{x}_{i,j}) = \frac{D}{V_i} \exp(-k_i(t_{ij} - t_D))$ where $\boldsymbol{\theta}_i = (V_i, k_i)$ and $\mathbf{x}_{i,j} = (D, t_{ij}, t_D)$.

To model the inter-patient variability, the individual parameters can be expressed as follows: $\boldsymbol{\theta}_i = \boldsymbol{\mu} \exp(\boldsymbol{\eta}_i)$, where $\boldsymbol{\eta}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega})$. $\boldsymbol{\mu}$ is the vector of fixed effects and represents the effect of a standard patient and $\boldsymbol{\Omega}$ is the variance covariance matrix of the random effects. More complex inter-individual models can be developed to account for potential covariates.

Finally, g represents the error model, where $\boldsymbol{\xi}$ are its parameters. Three error models are commonly used as defined below:

$$g(\boldsymbol{\theta}_i, \mathbf{x}_{i,j}, \boldsymbol{\xi}) = \begin{cases} a, & \boldsymbol{\xi} = a & (1) \\ bf(\boldsymbol{\theta}_i, \mathbf{x}_{i,j})^c, & \boldsymbol{\xi} = (b, c) & (2) \\ a + bf(\boldsymbol{\theta}_i, \mathbf{x}_{i,j})^c, & \boldsymbol{\xi} = (a, b, c) & (3) \end{cases} \quad (2.11)$$

where (1) corresponds to the constant error model, (2) corresponds to the proportional error model and (3) corresponds to the combined error model.

The notations were inspired from [45] and [5].

Parameters can be estimated by maximum likelihood. However, in case of non-linear mixed effects models, the likelihood does not have an analytical expression. Linearization of the likelihood has first been proposed [136] and parameters could be estimated with standard maximization algorithms (Newton-Raphson for example). But these methods only approximated the likelihood. Alternative methods that compute the “exact” likelihood using intensive calculation have then been developed such as the stochastic approximation expectation maximization (SAEM) algorithm [52] that is implemented in Monolix software [7]. More details can be found in Appendix B.

PK population modeling is illustrated in Figure 2.7 where the population response and the individual responses for each patient can be estimated from the

individual sampled data.




Figure 2.7 – Population PK modeling: population and individual profiles fitted from the individual sampled data represented by the dots.

Therefore, population PK/PD modeling aims to characterize the entire dose-exposure-response relationship through pharmacology principles in order to support the selection of the optimal dose regimen for the clinical development [142].

2.2 Dose-finding designs in oncology

As mentioned in the introduction, the objective of phase I trials is to determine the MTD. Many sequential dose-escalation designs have been proposed in the literature for various contexts, either considering a single administration or multiple administrations of the same or varying doses, and some designs have been proposed to extend the binary toxicity outcome, to include efficacy data or to consider the combination of two or more agents.

2.2.1 Single administration

2.2.1.1 Standard designs evaluating a binary toxicity outcome

Standard dose-finding designs have been developed to determine the MTD of the drug after a single administration (in the first cycle) using a binary toxicity outcome, that is the DLT. In the literature, these designs have been divided in three main classes that are the algorithm based designs, the model based designs and the model assisted designs.

2.2.1.1.1 Algorithm based designs Algorithm-based designs, also known as up & down designs, are simple dose-escalation designs that do not require any assumption on the dose-toxicity curve [100]. Dose allocation is based on algorithmic rules built on the number of observed DLTs.

The most widely known algorithm based design is the 3+3 design [144], that is illustrated in Figure 2.8.




Figure 2.8 – 3+3 design.

The trial begins at the lowest dose-level and dose-escalation is performed step by step. Three patients are included at the current dose-level. If none of the three

patients experience a DLT, dose-escalation is performed and three new patients are treated at the next higher dose-level. If at least two patients experience a DLT, the trial is stopped and the previous dose-level is recommended for the MTD. If only one patient experiences a DLT, three additional patients are treated with this current dose-level and if at least one of these new patients experiences a DLT, the trial is stopped; otherwise, dose-escalation is performed.

The main advantage of this design is its simplicity of implementation, but this simple design has been widely criticized. Indeed, the target of DLT probability is not specified and de-escalation is not possible, therefore a maximum of six patients are treated at each dose-level. Moreover, the design only uses the data from the most recent cohort (or the two most recent cohorts if one DLT is observed) to guide dose-escalation and therefore ignores the data from previous patients [42]. It has been shown that the 3+3 design has a lower performance in selecting the correct MTD compared to more complex designs [73] and has a tendency to recommend doses below the MTD.

Other algorithm based designs have been proposed to achieve better performance. Extensions of the 3+3 design, the so-called A+B designs [42], have been proposed to allow more flexibility in the cohort sizes and the number of DLTs to perform dose-escalation. These designs can also accommodate dose de-escalation to allow patients to be treated at lower doses in case of excessive toxicity. Other designs, defined as the group up-and-down designs, have been developed to target a prespecified target of toxicity [65, 75]. A review of various algorithm based designs is proposed in Liu et al. [100].

2.2.1.1.2 Model based designs Model based dose-escalation designs, where a parametric model is assumed for the dose-toxicity relationship, have been proposed to tackle the issues of simple algorithm based approaches. Most approaches have been developed in the Bayesian paradigm which is introduced in Appendix A.

The continual reassessment method (CRM) [119] has been developed to analyze all the information gathered at each step of the dose-escalation design with the aim of reducing the number of patients allocated to too toxic doses and increasing the number of patients treated at efficacious doses [62].

The general idea of the CRM is to fit a dose-toxicity model to the data in order to treat the patients with the dose most likely to have a probability of DLT close to the target. At the beginning of the trial, an initial dose-toxicity relationship is assumed. After each cohort of patients treated, the dose-toxicity model is updated with the entire data collected, and the next cohort of patients is treated with the estimated MTD that is the dose with the probability of toxicity closest to the target. The trial is stopped when the total number of patients is included.

The sequential process of the CRM is illustrated in Figure 2.9.

Let (d_1, \dots, d_K) denote the K dose-levels being investigated in the trial. Let $x_i \in \{d_1, \dots, d_K\}$ be the dose-level allocated to patient $i \in \{1, \dots, n\}$ and let $y_i \in \{0, 1\}$ be his/her toxicity response. Let $F(x_i, \boldsymbol{\theta})$ be the dose-toxicity model, that is a monotonic increasing function of the dose x_i and where $\boldsymbol{\theta} \in \Theta$ is the vector

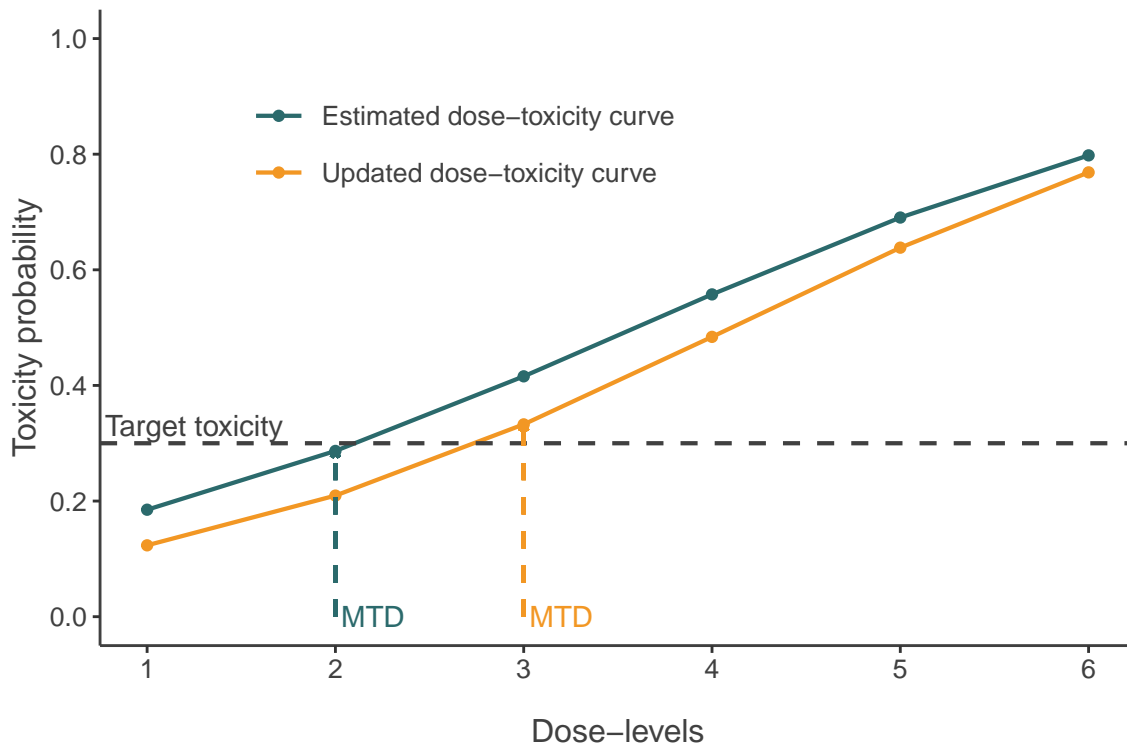


Figure 2.9 – Continual reassessment method (CRM) using a 1-parameter logistic model: the dose-toxicity curve is updated with the data of patients treated with the most recent estimated MTD (dose-level 2) and the MTD is updated (dose-level 3).

of parameters to estimate. After l patients are included, let $\mathbf{X}_l = (x_1, \dots, x_l)$ and $\mathbf{Y}_l = (y_1, \dots, y_l)$ be the data collected. The likelihood is defined as follows:

$$\mathcal{L}(\boldsymbol{\theta}; \mathbf{X}_l, \mathbf{Y}_l) = \prod_{i=1}^l F(x_i, \boldsymbol{\theta})^{y_i} (1 - F(x_i, \boldsymbol{\theta}))^{1-y_i} \quad (2.12)$$

Let g be the prior distribution of $\boldsymbol{\theta}$. The posterior distribution of $\boldsymbol{\theta}$ is defined as follows:

$$f(\boldsymbol{\theta} | \mathbf{X}_l, \mathbf{Y}_l) = \frac{\mathcal{L}(\boldsymbol{\theta}; \mathbf{X}_l, \mathbf{Y}_l) g(\boldsymbol{\theta})}{\int_{\Theta} \mathcal{L}(\boldsymbol{\theta}; \mathbf{X}_l, \mathbf{Y}_l) g(\boldsymbol{\theta}) d\boldsymbol{\theta}} \quad (2.13)$$

The posterior toxicity probability of each dose-level k can be estimated from the posterior mean of the parameters, defined as $\hat{\boldsymbol{\theta}} = \int_{\Theta} \boldsymbol{\theta} f(\boldsymbol{\theta} | \mathbf{X}_l, \mathbf{Y}_l) d\boldsymbol{\theta}$, as follows $\hat{\pi}_k = F(d_k, \hat{\boldsymbol{\theta}})$. The posterior toxicity probability can also be estimated as the posterior mean of the toxicity probability as follows $\hat{\pi}_k = \int_{\Theta} F(d_k, \boldsymbol{\theta}) f(\boldsymbol{\theta} | \mathbf{X}_l, \mathbf{Y}_l) d\boldsymbol{\theta}$.

The dose recommended for the next cohort of patients is estimated as d_{k^*} where $k^* = \min_{k \in \{1, \dots, K\}} |\hat{\pi}_k - \delta_T|$, where δ_T is the target probability of toxicity.

Various dose-toxicity models have been proposed, the main ones are defined below:

- Hyperbolic tangent model, where $\theta = a$: $F(x_i, \theta) = \left(\frac{\tanh(x_i) + 1}{2} \right)^a$, $a > 0$
- Power model, where $\theta = a$: $F(x_i, \theta) = x_i^{\exp(a)}$
- 1-parameter logistic model, where $\theta = b$: $F(x_i, \theta) = \frac{\exp(a_0 + \exp(b)x_i)}{1 + \exp(a_0 + \exp(b)x_i)}$, a_0 is fixed
- 2-parameter logistic model, where $\theta = (a, b)$: $F(x_i, \theta) = \frac{\exp(a + \exp(b)x_i)}{1 + \exp(a + \exp(b)x_i)}$

The dose-levels (d_1, \dots, d_K) can be the real doses administered (for example 1 mg/kg) or can be obtained from the initial guesses of the toxicity probabilities $(\pi_1^0, \dots, \pi_K^0)$, defined as the skeleton. In details, $d_k = F^{-1}(\pi_k^0, \theta_0)$, where θ_0 is the prior guess of the parameters value, for example the mean of the prior distribution [42]. For example, if θ follows a normal distribution centered in 0, we have $d_k = \pi_k^0$ for the power model and $d_k = \log\left(\frac{\pi_k^0}{1 - \pi_k^0}\right) - a_0$ for the 1-parameter logistic model. Lee and Cheung [89] proposed a method to define the initial guesses of toxicity probabilities to optimize the operating characteristics of the CRM when no prior information is available.

Various prior distributions can be considered, where a standard choice for the power model is $a \sim \mathcal{N}(0, 1.34)$ and for the 1-parameter logistic model with $a = 3$ is $b \sim \mathcal{N}(0, 1.34)$ [40]. Morita et al. [109] provided a definition of the prior effective sample size (ESS) in order to quantify the amount of information of the prior. A simple approximation of the prior ESS can be computed by matching the mean and variance of the toxicity probabilities computed from the prior distribution to those of a beta distribution [175].

Several extensions of the CRM, named modified CRM, have been proposed in the literature among which, the start of dose-escalation at the lowest dose, escalation to up to one dose-level, various cohort sizes or the inclusion of stopping rules other than the total sample size [42]. To avoid the number of patients receiving too toxic doses, the escalation with overdose control (EWOC) design has been developed [21] that assigns the dose having the posterior probability of exceeding the MTD equals a prespecified value (the feasibility bound). Neuenschwander et al. [116] have proposed a two-parameter model to have a more realistic representation of the dose-toxicity model. They also highlighted the limits of recommending the dose having the posterior toxicity probability closest to the target and instead divided the toxicity probability in four intervals (under-dosing, targeted toxicity, excessive toxicity and unacceptable toxicity) and proposed to recommend the MTD having the highest probability of toxicity in the targeted interval while controlling the probability of under and over dosing. Cheung and Chappell [41] have developed the time-to-event CRM (TITE-CRM) to include late onset toxicities while shortening trial duration. Indeed, in the TITE-CRM, the inclusion of a new cohort can be performed while previous patients may not be completely followed by considering a weighted dose-response model where the weight is a function of the time to event of the patient. Extensions of the TITE-CRM using a survival working model have

also been proposed [20].

Many authors have showed that model-based approaches exhibit better operating characteristics than traditional algorithm based approaches in terms of MTD selection and optimal dose allocation [73, 131, 84]. However these designs are more complex than algorithm-based designs and require computations after each cohort of patients to select the next dose to be administered. To allow the CRM to be more transparent to clinicians, Yap et al. [170] introduced the dose transition pathways that project the recommended doses for initial cohorts of patients in order for the clinicians to visualize the CRM decisions and provide feedback to further calibrate the design.

2.2.1.1.3 Model assisted designs Model-assisted designs have been developed to combine the simplicity of algorithm-based designs with the performance of model-based designs [169]. Dose escalation rules can be prespecified at the beginning of the trial similarly to algorithm based designs as illustrated in Figure 2.10, but these rules are computed from a statistical model [174]. However, contrary to model based designs, model assisted designs usually model only local data (data observed at each dose-level) with a binomial model without assuming a parametric model for the dose-toxicity relationship [182].

		Number of patients treated at current dose															
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
Number of dose limiting toxicities (DL T's)	0	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	
	1	D	S	S	S	S	E	E	E	E	E	E	E	E	E	E	
	2		DU	D	S	S	S	S	S	S	S	S	E	E	E	E	
	3			DU	DU	D	S	S	S	S	S	S	S	S	S	S	
	4				DU	DU	DU	D	D	S	S	S	S	S	S	S	
	5					DU	DU	DU	DU	DU	D	S	S	S	S	S	
	6						DU	DU	DU	DU	DU	DU	D	S	S	S	
	7							DU	DU	DU	DU	DU	DU	DU	D	S	
	8								DU	DU	DU	DU	DU	DU	DU	DU	DU
	9									DU	DU	DU	DU	DU	DU	DU	DU
	10										DU	DU	DU	DU	DU	DU	DU
	11											DU	DU	DU	DU	DU	DU
	12												DU	DU	DU	DU	DU
	13													DU	DU	DU	DU
	14														DU	DU	DU
	15															DU	DU

E = Escalate to the next higher dose
S = Stay at the current dose
D = De-escalate to the next lower dose
U = The current dose is unacceptably toxic
 MTD = 30%
 Sample Size = 30
 Epsilon1 = 0.05
 Epsilon2 = 0.05

Figure 2.10 – Decision rules for a model-assisted design (mTPI software downloaded from MD Anderson Cancer Center [6]).

Ji et al. [78] have proposed the mTPI method (calibration free extension of TPI [77]) where the partition $[0, 1]$ is divided in three intervals: low toxicity, acceptable toxicity (equivalence interval) and high toxicity. Decision rules to de-escalate, stay at the current dose or escalate are based on the unit probability mass of each interval (probability of the interval divided by the length of the interval calculated with a beta-binomial model) and additional safety rules. The mTPI-2 method [70] was introduced to correct undesirable decisions by defining multiple intervals of the same size. Yan et al. [169] have developed the keyboard design that has a better overdose control compared to the mTPI where equal width dosing intervals are defined (keys), and the decision is based on the strongest key having the highest posterior probability. Zhou et al. [182] showed that mTPI-2 and keyboard are equivalent, but mTPI-2 is less transparent due to the problem of Ockham’s razor. Liu and

Yuan [101] have proposed the Bayesian optimal interval design (BOIN) where decisions are made by comparing the observed toxicity rate with the interval boundaries that depend on the dose-level and the number of patients treated.

It has been shown that these model-assisted designs outperformed the 3+3 design and comparable results were found with model based approaches [182].

2.2.1.2 Designs evaluating an extension of the toxicity outcome

The designs presented in Section 2.2.1.1 have been developed for a binary toxicity endpoint, that is the DLT. However, these designs do not account for moderate toxicities (for example, adverse events of grade 2), and do not differentiate between two types of adverse events considered as a DLT (for example a grade 3 and a grade 4 adverse event or a grade 3 neuropathy and a grade 3 fatigue).

Wang et al. [161] extended the 3+3 design and the CRM to differentiate adverse events of grade 3 and grade 4. Yuan et al. [178] developed the quasi-CRM, an extension of the CRM for multiple toxicity grades by defined the equivalent toxicity score. Meter et al. [108] and Dousseau et al. [54] considered proportional odds models to include moderate toxicities in the estimation.

Bekele and Thall [23] assumed that various types of DLT might not be equally important and introduced a continuous variable, the total toxicity burden, that summarizes all toxicity events experienced by the patients using weights of the toxicity types and grades elicited by clinicians as illustrated in Table 2.1.

Type of toxicity	Grade	Severity weight
Myelosuppression without fever	3	1.0
	4	1.5
Myelosuppression with fever	3	5.0
	4	6.0
Dermatitis	3	2.5
	4	6.0
Liver	2	2.0
	3	3.0
	4	6.0
Nausea/vomiting	3	1.5
	4	2.0
Fatigue	3	0.5
	4	1.0

Table 2.1 – Toxicities and severity weights defined by Bekele and Thall [23].

Chen et al. [39] extended the equivalent toxicity score of Yuan et al. [178] to include multiple toxicities that can be correlated and various grades. Lee et al. [87] introduced the toxicity burden score, a weighted sum of grades and types of toxicities. Lee et al. [88] extended the CRM for a continuous or ordinal toxicity outcome (toxicity grades, total toxicity burden or toxicity burden score) to define the MTD under multiple toxicity constrains. Ezzalfani et al. [57] introduced the

total toxicity profile, and proposed to apply the quasi-CRM on the normalized total toxicity profile. Muenz et al. [111] proposed to model the total number of DLT and moderate toxicities. Lin [93] extended the BOIN designs to account for multiple toxicity grades and types using multiple toxicity constrains.

2.2.1.3 Designs evaluating efficacy in addition to toxicity

Phase I trials in oncology traditionally focus on toxicity that is considered as a surrogate endpoint for efficacy [42, 50] while phase II trials investigate efficacy. However, for some molecules, efficacy can reach a plateau at non toxic doses or can exhibit sufficient efficacy below the MTD [42]. Therefore, dose-finding designs that investigate both toxicity and efficacy have been proposed in phase I/II trials.

Ivanova [74] proposed to include efficacy data in the algorithmic rules. Many developments have been proposed for model-based approaches including efficacy data. Some authors considered a trinary outcome, no response (no toxicity and no efficacy), success (efficacy and no toxicity), and toxicity modeled by a proportional odds model or a continuation ratio model [151, 181]. Many authors developed bivariate models where toxicity and efficacy are explicitly modeled and the joint probability of toxicity and efficacy is expressed from the marginal distributions and an association parameter [113, 30, 148, 55, 171, 177].

In particular, Thall and Cook [148] and Dragalin and Federov [55] considered bivariate binary models and modeled the joint probability of efficacy and toxicity using an Archimedean copula [115] that has an explicit distribution defined from the marginal distributions of each variable and an association parameter. Let $\pi_T = \mathbb{P}(Y_T = 1)$ and $\pi_E = \mathbb{P}(Y_E = 1)$ denote the probability of toxicity and efficacy, respectively. An Archimedean copula is defined as follows:

$$C_\alpha(\pi_T, \pi_E) = \Phi_\alpha^{-1}(\Phi_\alpha(\pi_T) + \Phi_\alpha(\pi_E)), \quad (2.14)$$

where Φ_α is a continuous, strictly decreasing and convex function and Φ_α^{-1} is its inverse.

For example, the Farlie–Gumbel–Morgenstern distribution is defined as follows:

$$C_\alpha(\pi_T, \pi_E) = \pi_T \pi_E + \pi_T (1 - \pi_T) \pi_E (1 - \pi_E) \frac{\exp(\alpha) - 1}{\exp(\alpha) + 1}, \quad (2.15)$$

where $\alpha > 0$ for positive association and $\alpha \in [-1, 0[$ for negative association, and the Clayton distribution is defined as

$$C_\alpha(\pi_T, \pi_E) = (\max(\pi_T^{-\alpha} + \pi_E^{-\alpha} - 1, 0))^{-1/\alpha}, \quad (2.16)$$

where $\alpha > 0$ for positive association and $\alpha \in [-1, 0[$ for negative association.

The bivariate binary model is defined in terms of the four possible outcomes as follows [124]:

$$\begin{cases} \pi_{11} &= \mathbb{P}(Y_T = 1, Y_E = 1) = C_\alpha(\pi_T, \pi_E) \\ \pi_{01} &= \mathbb{P}(Y_T = 0, Y_E = 1) = \pi_E - C_\alpha(\pi_T, \pi_E) \\ \pi_{10} &= \mathbb{P}(Y_T = 1, Y_E = 0) = \pi_T - C_\alpha(\pi_T, \pi_E) \\ \pi_{00} &= \mathbb{P}(Y_T = 0, Y_E = 0) = 1 - \pi_T - \pi_E + C_\alpha(\pi_T, \pi_E) \end{cases} \quad (2.17)$$

However, Cunanan and Koopmeiners [51] showed that a simple model that assumes independence between the probability of toxicity and efficacy can perform as well as a copula model in the context of binary variables in phase I/II trials because the parameter of the copula model is difficult to be estimated. Koopmeiners and Modiano [79] proposed to jointly model the time to toxicity and the time to death with a copula model.

Bekele and Shen [114] considered a continuous biomarker outcome for efficacy and a probit model for toxicity with a latent variable. They modeled the joint distribution with a bivariate normal distribution where the covariance parameter could model the association between the biomarker and toxicity.

Finally, some authors modeled the probability of efficacy conditionally on no toxicity [117, 168, 183, 32]. Extensions of model-assisted designs with efficacy have also been developed [91, 96, 146, 97].

2.2.1.4 Designs evaluating drug combinations

The dose-finding designs presented above assume that a single agent is being evaluated. Administering a combination of two (or more) agents, that have already been investigated in previous trials, can increase potential anti-tumor activity due to a synergistic effect in terms of efficacy. However, a synergistic effect in terms of toxicity can also be observed. Therefore, dose-finding combination trials are performed to find the maximum tolerated drug combination. Contrary to single agent trials, a complete order in terms of toxicity cannot be assumed, as illustrated in Figure 2.11. Toxicity probability is assumed to increase when increasing the dose of one agent, while keeping the second agent fixed. However, some combinations cannot be initially ordered, for example in Figure 2.11, we do not know at the beginning of the trial if drug combination d_{12} is more toxic than drug combination d_{21} .

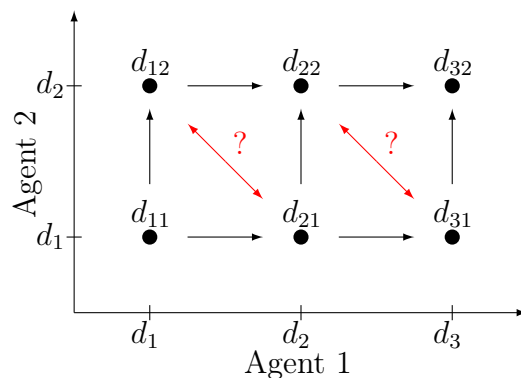


Figure 2.11 – Partial order in the combination of two agents.

The simplest design for drug combination trials is to define a series of subtrials where one agent is varied for each trial [76] to ignore the partial order, but this design requires a large sample size. Conaway et al. [48] proposed a non parametric design where the treatment is randomly selected in case of multiple choices. Fan et al. [58] developed an algorithm-based design with isotonic regressions where two tied treatments can be tested simultaneously.

Yuan and Yin [176] divided the 2 dimensional trial in groups of subtrials and conducted the subtrials sequentially with a CRM where the information from each subtrial is used to shrink the information for remaining subtrials, it was extended by Zhang and Yuan [180]. Wages et al. [159] developed the partial order CRM (POCRM) where a subset of possible orderings is prespecified for the combinations [158] and a CRM is run on each complete ordering. The next combination is recommended for the order having the highest posterior probability. Braun and Jia [31] developed a generalized version of the CRM for combination trials where the intercept of the logistic model is a parameter specific to the second agent.

Other authors explicitly modeled the probability of toxicity of the drug combination with the dose of each agent [149, 162, 104, 22, 37, 130]. Yin and Yuan [172, 173] modeled the probability of toxicity of the combination from the probability of toxicity of each agent using a copula model similarly to toxicity and efficacy models introduced in Section 2.2.1.3. A comparison of some algorithm based and model based dose-finding designs for combination trials can be found in [129].

Extensions of model-assisted designs for drug combinations have also been proposed [95, 103, 120].

2.2.2 Multiple administrations

The dose-finding designs presented in Section 2.2.1 have been developed to find the MTD after a single administration of the drug (or combination), although anti-cancer therapies are usually administered in multiple cycles. Some designs have been developed to account for multiple administrations in order to optimize the schedule of administration or evaluate the sequence of the same or varying doses.

2.2.2.1 Dose-schedule finding as a two-dimensional problem

Dose-schedule finding can be seen as a two-dimensional problem, similarly to the combination of two agents, where the dose is the first dimension and the schedule is the second. However, while the probability of toxicity is assumed to increase with each agent for combinations, the probability of toxicity might not be ordered with the schedule.

For example, let's consider two cases where three per-administration dose and two schedules are under evaluation.

For Case 1, the two evaluated schedules are defined as:

- Schedule A: Administration on days 1-3 of a 21 days cycle
- Schedule B: Administration on days 1-6 of a 21 days cycle

For Case 2, the two evaluated schedules are defined as:

- Schedule A: Administration on days 1-3 and 8-10 of a 21 days cycle
- Schedule B: Administration on days 1-6 of a 21 days cycle

For Case 1, schedule B is assumed to be more toxic than schedule A, therefore we have the same orderings than for combination trials as illustrated in the left of Figure 2.12. However, for Case 2, schedules A and B cannot be initially ordered as illustrated in the right of Figure 2.12.




Figure 2.12 – Partial order in dose-schedule finding as a two-dimensional problem.

A simple approach consists in evaluating the MTD for each schedule, however this approach requires a large sample size when multiple schedules are investigated.

Some authors have developed dose-schedule finding designs in this two-dimensional framework, assuming the schedules can be ordered. The design proposed by Yuan and Yin [176], presented in Section 2.2.1.4, also applied to dose-schedule finding when the schedules can be ordered. Li et al. [92] developed a dose-schedule finding design using both efficacy and toxicity using a global cross ratio model and proposed a Bayesian isotonic transformation to preserve the partial ordering. Quintana et al. [126] proposed a Bayesian design to find the optimal dose and schedule combination evaluating toxicity and efficacy in multiple disease subtypes with logistic models in the context of two doses and two ordered schedules under study.

Other authors have considered the case where the schedules cannot be initially ordered. Thall et al. [150] developed a 2-stage design to optimize the dose and schedule based on the joint utility of the times to toxicity and response where in stage 1, patients are randomized among schedules. Wages et al. [160] also relaxed the assumption of complete ordered schedules and extended the POCRM (developed for combination trials) for dose-schedule finding where various subsets of complete orderings are prespecified. Guo et al. [69] proposed a dose-schedule finding design for non ordered schedules based on a trinary outcome to account for toxicity and efficacy using a Bayesian dynamic model. They developed a multistage design where patients are randomized among schedules for the first step. Lin et al. [94] proposed a Bayesian hierarchical model to optimize the dose-schedule combination using utilities to quantify the trade-off between efficacy and toxicity.

These designs consider the schedules as qualitative and therefore do not include the actual timing of drug administrations in the modeling. Moreover, they assume that the same dose is administered repeatedly.

2.2.2.2 Modeling the time to DLT after multiple administrations

Another definition of schedule has emerged in the literature as the duration of treatment where the objective is to find the optimal number of courses of therapy accounting for cumulative effect of repeated administrations.

Braun et al. [33] extended the TITE-CRM where the schedule is considered as the “dose” in order to determine the maximum tolerated cumulative dose. However, this method does not account for incomplete schedules and the contribution of overlapping “doses” to DLT was unclear. These findings motivated the modeling of how toxicity is related to each administration during treatment.

Braun and Yuan [35] proposed to model the time to toxicity considering each administration of a same dose to determine the maximum tolerated schedule, where the schedule denotes the number of courses of therapy. They defined the total hazard of toxicity of a patient at study time t as the sum of hazards related to each administration as follows:

$$\lambda(t|\boldsymbol{\theta}, \mathbf{s}) = \sum_{l=1}^m h(t - s_l|\boldsymbol{\theta}), \quad (2.18)$$

where \mathbf{s} is the schedule (until t) defined as the sequence of m administrations at times $\mathbf{s} = (s_1, \dots, s_m)$ and $\boldsymbol{\theta}$ is the vector of the model parameters.

They defined the survival function as follows:

$$S(t|\boldsymbol{\theta}, \mathbf{s}) = \mathbb{P}(Y > t|\boldsymbol{\theta}, \mathbf{s}) = \exp\left(-\sum_{l=1}^m H(t - s_l|\boldsymbol{\theta})\right) \quad (2.19)$$

They proposed a 3-parameter triangular function, $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$, for the single hazard function, illustrated in Figure 2.13, where the hazard increases and reaches its maximum θ_2 at time θ_1 and then decreases to 0 at time θ_3 as follows:

$$h(u|\boldsymbol{\theta}) = \begin{cases} \theta_2 \frac{u}{\theta_1}, & 0 \leq u \leq \theta_1 \\ \theta_2 \frac{\theta_3 - u}{\theta_3 - \theta_1}, & 0 < u \leq \theta_3 \\ 0, & u > \theta_3 \text{ or } u < 0 \end{cases} \quad (2.20)$$

Braun et al. [34] extended this previous work to determine a maximum tolerated dose and schedule to allow the per-administration dose to vary in addition to the schedule. They reparameterized the triangular function and defined one single hazard per dose d_j as follows:

$$h(u|\boldsymbol{\theta}_j) = \begin{cases} \frac{2a_j}{b_j + c_j} \frac{u}{b_j}, & 0 \leq u \leq b_j \\ \frac{2a_j}{b_j + c_j} \frac{b_j + c_j - u}{c_j}, & b_j < u \leq b_j + c_j \\ 0, & u > b_j + c_j \text{ or } u < 0 \end{cases} \quad (2.21)$$

where $a_1 < a_2 < \dots < a_j$ for the cumulative hazard toxicity to increase with the dose.




Figure 2.13 – Triangular and Weibull hazard functions.

Liu and Braun [99] proposed a parametric non-mixture cure model to identify the maximum tolerated schedule with a fixed per-administration dose. Cured patients are defined as patients who do not experience DLT and the authors defined the hazard of each administration as being proportional to a 2-parameter Weibull density function, illustrated in Figure 2.13, and defined as follows:

$$f(t|\boldsymbol{\phi}) = \exp(-\gamma) \alpha t^{\alpha-1} \exp(-t^\alpha \exp(-\gamma)) \quad (2.22)$$

The survival function is defined as follows:

$$S(t|\boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}) = \mathbb{P}(Y > t|\boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}) = \exp\left(-\theta(m|\boldsymbol{\beta}) \sum_{l=1}^m F(t - s_l|\boldsymbol{\phi})/m\right), \quad (2.23)$$

where $\log(\theta(m|\boldsymbol{\beta})) = \beta_0 + \beta_1 \log(m)$, $\beta_1 > 0$, for the cumulative probability of DLT to increase with the number of administrations.

Zhang and Braun [179] extended the previous work to optimize the dose in addition to the schedule both between and within patients. They also considered a Weibull density but defined a dose dependent cure rate parameter. The survival function is defined as follows:

$$S(t|\boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}, \mathbf{d}) = \mathbb{P}(Y > t|\boldsymbol{\beta}, \boldsymbol{\phi}, \mathbf{s}, \mathbf{d}) = \exp\left(-\sum_{l=1}^m \theta_l F(t - s_l|\boldsymbol{\phi})\right), \quad (2.24)$$

where $\log(\theta_l) = \beta_0 + \exp(\beta_1) d_l$.

These approaches account for cumulative effects of multiple administrations in the hazard function but assumes uncorrelation between the administrations. However, for some molecules, the effect of some administrations can be associated with the previous ones.

2.2.2.3 Evaluating the sequence of administrations

Other authors evaluated multiple administrations of the drug as a sequence of the same or varying doses.

Simon et al. [139] investigated the cumulative effect of multiple courses of treatment on the risk of DLT and developed the accelerated titration designs. They are algorithm-based designs developed to shorten trial duration, treat fewer patients at sub therapeutic doses and integrate cumulative toxicities, that can allow inpatient dose-escalation.

Some authors evaluated the effect of a same dose being given in multiple cycles. Huang and Kuan [72] extended the TITE-CRM for multiple cycles of treatment using an adaptive weight function to include cycle information.

Dousseau et al. [54] proposed a mixed-effect proportional odds model to account for the longitudinal measurements of toxicity in multiple cycles. They assumed that the same dose was given in subsequent cycles and determined the RP2D from the estimated probabilities of DLT per cycle. They also proposed to detect a time trend at the end of the trial. Paoletti et al. [121] showed that using the information on multiple cycles increases the accuracy to determine the MTD.

Ezzalfani et al. [56] extended the design developed by Bekele and Shen [114] to model the repeated toxicity and biomarker data for two cycles of treatment. They used linear models considering a latent variable for toxicity and recommended the dose based on the probability of DLT at each cycle and the efficacy probability.

Altzerinakou and Paoletti [19] proposed to jointly model a continuous biomarker and time-to-DLT data for multiple cycles of the same dose using a shared random effect. Dose recommendation was based on efficacious doses that do not exceed a target cumulative probability of DLT over the treatment cycles.

Fernandes et al. [59] proposed to model the conditional probability of toxicity on any cycle given that no toxicity was observed in the previous cycles. Their Markov model is able to recommend a dose for future cycles, allowing inpatient dose escalation and deescalation, and to recommend an entire sequence of doses at the end of the trial. They defined the conditional probability of toxicity at cycle k , given no toxicity was observed before, as follows:

$$p_{i,k}^{\text{cond}} = \mathbb{P}(Y_{i,k} = 1 | Y_{i,k-1} = 0, \dots, Y_{i,1} = 0) \quad (2.25)$$

They proposed to model this probability of toxicity with the current dose administered, $d_{i,k}$, the maximum dose administered in the previous cycle, $d_{i,k-1}^{\text{max}} = \max_{j \in \{1, \dots, k-1\}}(d_{i,j})$, and the cumulative dose, $D_{i,k-1} = \sum_{j=1}^{k-1} d_{i,j}$ where $D_{i,1} = 0$, as follows:

$$\log(1 - p_{i,k}^{\text{cond}}) = -\alpha(d_{i,k} - \rho d_{i,k-1}^{\text{max}}) \mathbb{1}_{\{d_{i,k} > \rho d_{i,k-1}^{\text{max}}\}} - \beta D_{i,k} d_{i,k} \quad (2.26)$$

In the previous equation, ρ reflects the memory of non toxic previous doses, meaning that if a patient does not experience toxicity at a given dose, then he/she

might have a lower risk of toxicity at a higher dose. β reflects the potential cumulative toxicities that are the remaining damage caused by previous doses.

Ursino et al. [154] defined the cumulative probability of toxicity after a sequence of doses $(d_{i,1}, \dots, d_{i,k})$ defined as

$$p_{i,k}^{\text{cum}} = \mathbb{P}(\exists l \in \{1, \dots, k\} Y_{i,l} = 1) \quad (2.27)$$

They proposed to model this cumulative probability of toxicity with the first dose administered and the cumulative dose, defined as $D_{i,k} = \sum_{l=2}^k d_{i,k}$ where $D_{i,1} = 0$, as follows:

$$p_{i,k}^{\text{cum}} = \frac{\exp\left(\alpha + \exp(\beta) \log\left(\frac{d_{i,1}}{d^*}\right) + \exp(\gamma) \log\left(\frac{D_{i,k}}{D^*} + 1\right)\right)}{1 + \exp\left(\alpha + \exp(\beta) \log\left(\frac{d_{i,1}}{d^*}\right) + \exp(\gamma) \log\left(\frac{D_{i,k}}{D^*} + 1\right)\right)} \quad (2.28)$$

Lee et al. [86] developed a within and between patients adaptive design to dynamically optimize the dose of two cycles of therapy, similarly to dynamic treatment regimes [112]. They proposed a Bayesian hierarchical model for binary toxicity and efficacy where decisions are based on per-cycle utilities. Lyu et al. [102] developed a simple Bayesian dose-cycle adaptive design to identify the maximum tolerated sequence that allows inpatient modifications. Decisions are based on a beta binomial model and they developed a dose continuation rule to decide if the patients should receive another cycle of therapy and a dose allocation rule to select the dose, similarly to the mTPI-2 design. They proposed a method based on the isotonic transformation of the estimated probability of toxicity per cycle for each sequence to determine the maximum tolerated sequence.

Many dose-finding designs have been developed in various contexts including the evaluation of multiple administrations of the drug where the objective is to optimize the schedule of administration, the timing or the dose of each administration. Some designs account for the cumulative effect of multiple administrations, mainly by the cumulative dose or by summing the hazard related to each administration. However, the effect of more complex scheme of administration could be related to principles of pharmacology.

2.3 Dose-finding designs incorporating pharmacokinetics

Characterizing the PK profile of the drug is usually a secondary endpoint in phase I trials and can support the recommendation of the RP2D, while PD is usually an exploratory endpoint for efficacy. For example, PK studies evaluating different dosing schedules can help to determine the optimal route or dosing frequency [49]. However, in a survey conducted by Comets and Zohar [46], the PK study was usually described separately from the main findings of phase I trials in oncology. They also suggested to attempt to relate PK with the main results of the trial, such as toxicity. In the literature, some dose-finding designs that include PK have been developed.

2.3.1 Single administration

When designing a dose-finding trial, the starting dose can be defined from toxic doses observed in preclinical trials. However, Collins et al. [44, 43] argued that the comparison between species based on the dose may be inaccurate and may lead to a high number of dose-escalation steps. They proposed the algorithm-based pharmacologically guided dose escalation design where dose-escalation is guided by the measured AUC with the objective to reach a prespecified value, obtained from preclinical trials. Mao and Cheung [105] developed an individualized dosing algorithm to treat patients according to a prespecified biomarker or pharmacokinetic value.

Piantadosi and Liu [125] proposed to add the AUC as a covariate in the toxicity model. Patterson et al. [123] and Whitehead et al. [166, 167] proposed to model the AUC with the dose for phase I trials of healthy volunteers. Ursino et al. [155] extended these approaches to propose various dose-finding designs including PK:

- PKCOV (modification of Piantadosi and Liu [125]): The AUC is included as a covariate in the toxicity model. Dose-escalation is based on the target probability of toxicity.
- PKLIM (modification of Patterson et al. [123] and Whitehead et al. [166]): The AUC is modeled with the dose using a hierarchical model. Dose-escalation is based on a threshold on the AUC.
- PKCRM (combination of PKLIM and CRM): Dose-escalation is based on the lowest dose recommended by PKLIM and CRM.
- PKLOGIT (inspired by Whitehead et al. [167]): The AUC is modeled with the dose using a hierarchical model and the probability of toxicity is modeled with the AUC using a logit model. Dose-escalation is based on the target probability of toxicity. The approach is detailed below.
- PKPOP: Variation of PKLOGIT where the probability of toxicity is modeled with the mean AUC predicted by the hierarchical model.
- PKTOX: Modification of PKLOGIT with a probit model for the probability of toxicity.

For PKLOGIT, they assumed linear kinetics meaning that the AUC is proportional to the dose. They proposed to model the AUC to include inter-patient variability as follows:

$$\begin{cases} AUC_i | \boldsymbol{\beta}, \nu & \sim \mathcal{N}(\beta_0 + \beta_1 \log(d_i), \nu^2) \\ \boldsymbol{\beta} | \nu & \sim \mathcal{N}(\mathbf{m}, \nu^2 D_2) \\ \nu & \sim \text{Beta}(a, b) \end{cases} \quad (2.29)$$

where D_2 is a diagonal matrix and \mathbf{m} , D_2 , a , b are elicited values.

They modeled the probability of toxicity as follows:

$$\mathbb{P}(Y_i = 1 | AUC_i, \boldsymbol{\theta}) = \frac{\exp(-\theta_0 + \theta_1 AUC_i)}{1 + \exp(-\theta_0 + \theta_1 AUC_i)} \quad (2.30)$$

The next dose was then chosen as the one having the probability of toxicity closest to the target, where the probability of toxicity was defined as follows:

$$\mathbb{P}(Y = 1 | d, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\beta}}, \hat{\nu}) = \int_{\mathbb{R}} \frac{\exp(-\hat{\theta}_0 + \hat{\theta}_1 z)}{1 + \exp(-\hat{\theta}_0 + \hat{\theta}_1 z)} g(z | d, \hat{\beta}_1, \hat{\beta}_2, \hat{\nu}) dz, \quad (2.31)$$

where

$$g(z | d, \hat{\beta}_1, \hat{\beta}_2, \hat{\nu}) = \frac{1}{\hat{\nu} \sqrt{2\pi}} \exp\left(-\frac{1}{2} \left(\frac{z - (\hat{\beta}_0 + \hat{\beta}_1 \log(d))}{\hat{\nu}}\right)^2\right) \quad (2.32)$$

They compared these methods with the CRM and found that the methods that model both the dose-AUC and AUC-toxicity relationships (PKTOX and PKLOGIT) can provide a more precise estimation of the entire dose-toxicity curve without damaging the determination of the MTD.

Taketa et al. [145] proposed a modified version of PKLOGIT based on a standardized adjustment of the AUC.

2.3.2 Multiple administrations

Legedza and Ibrahim [90] developed a dose-response model based on PK principles to account for cumulative toxicities due to multiple dose administrations. They proposed to model the probability of toxicity with the total amount of drug in the bloodstream, that includes the current and the previous doses and accounts for the process of elimination of the drug that is linked with the clearance parameter (assuming an underlying 1 compartment bolus model). However, the elimination parameter could not be estimated with binary data and they proposed to elicit this parameter from an expert. They also proposed to account for the interpatient variability by adding a random effect.

Günhan et al. [67] developed a Bayesian time-to-event pharmacokinetic (TITE-PK) model to determine the maximum tolerated dose-schedule combination using PK. They defined the hazard of DLT as a function proportional to an exposure measure of the drug as follows:

$$h(t) = \beta E(t) \quad (2.33)$$

Drug exposure was related to drug concentration, however as measures of drug concentration are usually not directly available, they considered PK as a latent variable. They assumed a 1-compartment model and considered the drug concentration in the effect compartment to delay the effect from drug administration.

Let a same dose d be administered at a dosing frequency f , with a volume of distribution of 1 for simplicity, let k be the elimination parameter and let k_e be the equilibrium rate constant between the central and effect compartment. The concentration in the effect compartment was defined as follows:

$$C_e(t|d, f) = d \sum_{i=0}^{\inf} \mathbb{1}_{\{t \geq \frac{i}{f}\}} \frac{k_e}{k_e - k} \left(\exp\left(-k\left(t - \frac{i}{f}\right)\right) - \exp\left(-k_e\left(t - \frac{i}{f}\right)\right) \right), \quad (2.34)$$

where k and k_e were assumed to be known at the beginning of the trial, from preclinical studies for example.

The exposure measure was obtained from a reference dose-schedule combination (d^*, f^*) at the end of cycle 1 (t^*) as follows:

$$E(t|d, f) = \frac{C_e(t|d, f)}{\int_0^{t^*} C_e(t|d^*, f^*) dt} \quad (2.35)$$

The probability of DLT at the end of cycle 1 was then defined as follows:

$$\mathbb{P}(T \leq t^*|d, f) = 1 - \exp(-\beta AUC_E(t^*)), \quad (2.36)$$

where $AUC_E(t^*) = \int_0^{t^*} E(t|d, f) dt$.

Dose-escalation was based on an adapted EWOC criterion on the probability of DLT at the end of cycle 1 as follows:

$$\mathbb{P}(\mathbb{P}(T \leq t^*|d, f) > 0.33) < a, \quad (2.37)$$

where a is the feasibility bound and among the combinations that fulfill this criterion, the combination with the lowest $AUC_E(t^*)$ was chosen.

TITE-PK was shown to have good performance in finding a suitable dose-schedule combination compared to POCRM using PK to combine the information from different treatment schedules. Günhan et al. [68] adapted the TITE-PK model in trials where the schedules are investigated sequentially for the information on the different schedules to be directly considered through the PK principles. However, TITE-PK assumes that the probability of DLT at the end of cycle 1 is proportional to the measure of exposure. In the context of linear kinetics, this means that the probability of DLT at the end of cycle 1 is also proportional to the dose.

Few dose-finding designs including PK have been proposed in the literature and they usually rely on a simplified PK model. The main difficulty from including a complete PK (or PK/PD) model in the sequential dose-allocation design is that

PK/PD samples are usually not directly available, and first PK/PD model are usually developed after several patients have been treated. However, relating PK or PK/PD to toxicity and characterizing the entire dose-exposure response could improve the choice of the dose or dose regimen for future phases of the clinical development and could account for complex assumptions on the probability of toxicity, such as violation of the monotonicity assumption.

Chapter 3

Objectives of the thesis

3.1 Motivating trial

This thesis was motivated by a phase I/II first in human dose-escalation trial of SAR440234 in patients with relapsed or refractory acute myeloid leukemia, high risk myelodysplastic syndrome, or B cell acute lymphoblastic leukemia (ClinicalTrials.gov Identifier NCT03594955 [14]).

SAR440234 is a novel anti-CD123/CD3 bispecific T cell engager (BiTE) antibody that binds to both CD3 expressed on T cells and CD123 expressed on tumor cells, as shown in Figure 3.1. It activates and redirects cytotoxic T cells to CD123 expressing tumor cells leading to enhanced cytotoxic T cells mediated elimination of CD123-expressing tumor cells [15].




Figure 3.1 – Schematic representation of SAR440234, inspired by [106]

Cytokine release syndrome (CRS), a systemic inflammatory response, is one of the most frequent adverse events observed in T cell-engaging immunotherapies, that include BiTE [137].

Different symptoms can be associated with various severity of CRS, many of them mimic infections (fever for example) [85]. The NCI CTCAE [8] provides the

classification of CRS in the five toxicity grades, as shown in Table 1.1 in Chapter 1. However, Lee et al. [85] proposed a revised grading system that has been widely adopted and is defined in Table 3.1.

Grade 1	Symptoms are not life threatening and require symptomatic treatment only, eg, fever, nausea, fatigue, headache, myalgias, malaise
Grade 2	Symptoms require and respond to moderate intervention Oxygen requirement $<40\%$ or Hypotension responsive to fluids or low dose of one vasopressor or Grade 2 organ toxicity
Grade 3	Symptoms require and respond to aggressive intervention Oxygen requirement $\geq 40\%$ or Hypotension requiring high dose or multiple vasopressors or Grade 3 organ toxicity or grade 4 transaminitis
Grade 4	Life-threatening symptoms Requirement for ventilator support or Grade 4 organ toxicity (excluding transaminitis)
Grade 5	Death

Table 3.1 – Revised CRS grading from Lee et al. [85].

Cytotoxic T cell activation leads to the massive release of inflammatory cytokines [85] as shown in Figure 3.1. According to the NCI, cytokine is a “type of protein that is made by certain immune and non-immune cells and has an effect on the immune system” [10]. Interleukins (IL) and interferons (IFN) are examples of cytokines. Teachey et al. [147] found that the peak levels of many cytokines, such as IL-6, IL-10, and INF- γ , were associated with severe CRS. For Shimabukuro-Vornhagen et al. [137], IL-6 seems to have a key role in CRS.

Giving repeated doses of the drug can reduce the risk of CRS due to the mitigation of the peak of cytokine [38]. Priming dose strategies, where initial lower doses are administered followed by higher maintenance doses, have already been proposed to reduce the risk of severe CRS [38, 140, 153, 152, 157, 1].

The motivating trial is therefore designed with a fixed intra-patient dose-escalation scheme, where a dose regimen can be defined as a sequence of increasing doses, the lead-in doses, administered until the steady-state dose, which is then given repeatedly. An example of dose regimens with intra-patient dose-escalation is provided in Table 3.2.

	Day 1	Day 5	Day 9	Day 13	Day 17	Day 21	Day 25
S_1	1	5	10	10	10	10	10
S_2	1	5	10	20	20	20	20
S_3	5	10	20	40	40	40	40
S_4	5	10	20	60	60	60	60
S_5	10	20	60	100	100	100	100
S_6	10	20	60	120	120	120	120

Table 3.2 – Example of a panel of dose regimens with intra-patient dose-escalation (doses in $\mu\text{g}/\text{kg}$).

In this example, each dose regimen is composed of seven drug administrations (on days 1, 5, 9, 13, 17, 21 and 25) of increasing lead-in doses until the repeated steady-state dose. For instance, for dose regimen \mathbf{S}_4 , the first three drug administrations are the increasing lead-in doses, and the final four administrations are the repeated administrations of the steady-state dose ($60 \mu\text{g}/\text{kg}$). In this example, we assume a complete order of the DLT probabilities meaning that S_1 is the less toxic dose regimen and S_6 is the most toxic one due to a more aggressive lead-in doses scheme and a higher value of the steady-state dose when increasing the level of the dose regimen.

In addition to CRS, other toxicities can occur and the association between CRS and other toxicities has been investigated in the literature. Few works investigating Chimeric antigen receptor T cell therapies, another type of T cell-engaging immunotherapies, showed that CRS and neurotoxicity might be correlated. Wang and Han [163] argued that, while the exact mechanism generating neurotoxicity is not completely understood, CRS and neurotoxicity should not be considered as completely unrelated. Santomasso et al. [132] found a significant correlation of neurotoxicity with the presence and severity of CRS. According to Siegler and Kenderian [138], immune effector cell-associated neurotoxicity syndrome often accompanies and correlates with CRS, but it has also been occasionally reported to occur independently from CRS.

Further, other toxicities were found to be correlated to CRS, as hematologic toxicities [60]. Finally, Wei et al [165] stated that CRS can be affected by many factors including tumor burden, individual immune status or IL-6 levels and severe CRS can lead to multiple organ dysfunctions. In the following, we name DLT_o the DLT different from the CRS, that can for example include neurotoxicities.

The primary outcome of the trial is the incidence of DLT defined using the NCI CTCAE for DLT_o and the consensus guidelines for the CRS defined by Lee et al. [14, 85]. The primary objective of the trial is to find the MTD of the drug and the RP2D for the expansion part of the trial [14]. The 3+3 design, described in Section 2.2.1.1, was chosen for the dose-escalation design [27].

However, standard dose-escalation designs, such as the 3+3 or the CRM, ignore the intra-patient dose-escalation information and consider each dose regimen as a single “dose-level”. Table 3.3 illustrates an example of the 3+3 design applied on dose regimens with fixed intra-patient dose-escalation.

	Day 1	Day 5	Day 9	Day 13	Day 17	Day 21	Day 25	DLT/Patients
\mathbf{S}_1	1	5	10	10	10	10	10	0/3
\mathbf{S}_2	1	5	10	20	20	20	20	0/3
\mathbf{S}_3	5	10	20*	40	40	40	40	1/6
\mathbf{S}_4	5*	10	20*	60	60	60*	60	3/6
\mathbf{S}_5	10	20	60	100	100	100	100	0/0
\mathbf{S}_6	10	20	60	120	120	120	120	0/0

Table 3.3 – Example of a 3+3 design applied on the panel of dose regimens defined with intra-patient dose-escalation. The star represents the occurrence of a DLT.

In this example, no DLT occurs when administering dose regimens \mathbf{S}_1 and \mathbf{S}_2 . One DLT occurs when including 3 patients at dose regimen \mathbf{S}_3 , but no additional

DLT is observed when including 3 additional patients, so dose-escalation can continue. One DLT occurs when including 3 patients at dose regimen \mathbf{S}_4 , and two additional DLTs are observed when including 3 additional patients. The trial is therefore stopped, and the 3+3 design recommends dose regimen \mathbf{S}_3 as the MTD. However, when looking at the timing of DLT occurrences, two of three DLTs observed at dose regimen \mathbf{S}_4 occur during the lead-in doses (administrations on days 1 and 9), and \mathbf{S}_3 and \mathbf{S}_4 share the same lead-in doses. Therefore, one can wonder if \mathbf{S}_3 is safe enough to be recommended as the MTD.

Of note, a CRM would also reduce the dose regimen to a single dose-level since the method relies on a working model defined from the initial guesses of the DLT probabilities. A modified CRM could rely on a numeric dose and the dose regimen could then be simplified considering only the final dose planned (the steady state dose of the dose regimen).

The cytokine response can be seen as a PD endpoint. However, as the motivating trial was still ongoing at the time of my thesis, the PK/PD models for drug concentration and the cytokine response were not yet developed.

PK/PD models for drug concentration and cytokine response have been published for blinatumomab, a BiTE that binds to both CD3 expressed on T cells and CD19 expressed on B cells [2].

Zhu et al. [184] selected a one compartment linear model for the concentration of blinatumomab administered in continuous IV. They also found that the cytokines values increased quickly with the dose after the infusion of blinatumomab during the first week of therapy, among which IL-10, IL-6, and IFN- γ .

Chen et al. [38] developed a PD model for the cytokine response of blinatumomab from a phase I trial in patients with relapsed non-Hodgkin lymphoma. To model cytokine mitigation with multiple administrations, they assumed that cytokine production was stimulated by the release, RL, but inhibited by the negative feedback, IH, as illustrated in Figure 3.2.




Figure 3.2 – Cytokine model from Chen et al. [38].

For our work, we considered the following PD model where the release RL is linked with drug concentration C and the negative feedback IH is linked with the AUC of the PD response, AUC_R ,

$$\frac{dR(t)}{dt} = \frac{E_{\max}C(t)^H}{EC_{50}^H + C(t)^H} \left(1 - \frac{I_{\max}AUC_R(t)}{\frac{IC_{50}}{K^{N-1}} + AUC_R(t)} \right) - k_{\text{deg}}R(t), \quad (3.1)$$

where E_{\max} is the maximum cytokine release rate, EC_{50} is drug concentration to achieve 50% of maximum cytokine release rate, H is the Hill coefficient, I_{\max} is the maximum inhibition, IC_{50} is the cumulative cytokine exposure to achieve 50%

of the maximum inhibition, K is the priming factor and N is the number of doses. Values used for the PK/PD parameters can be found in Appendix B.1.

These PK/PD models are able to describe the reduction of the peak of cytokine with intra-patient dose-escalation as illustrated in Figure 3.3.

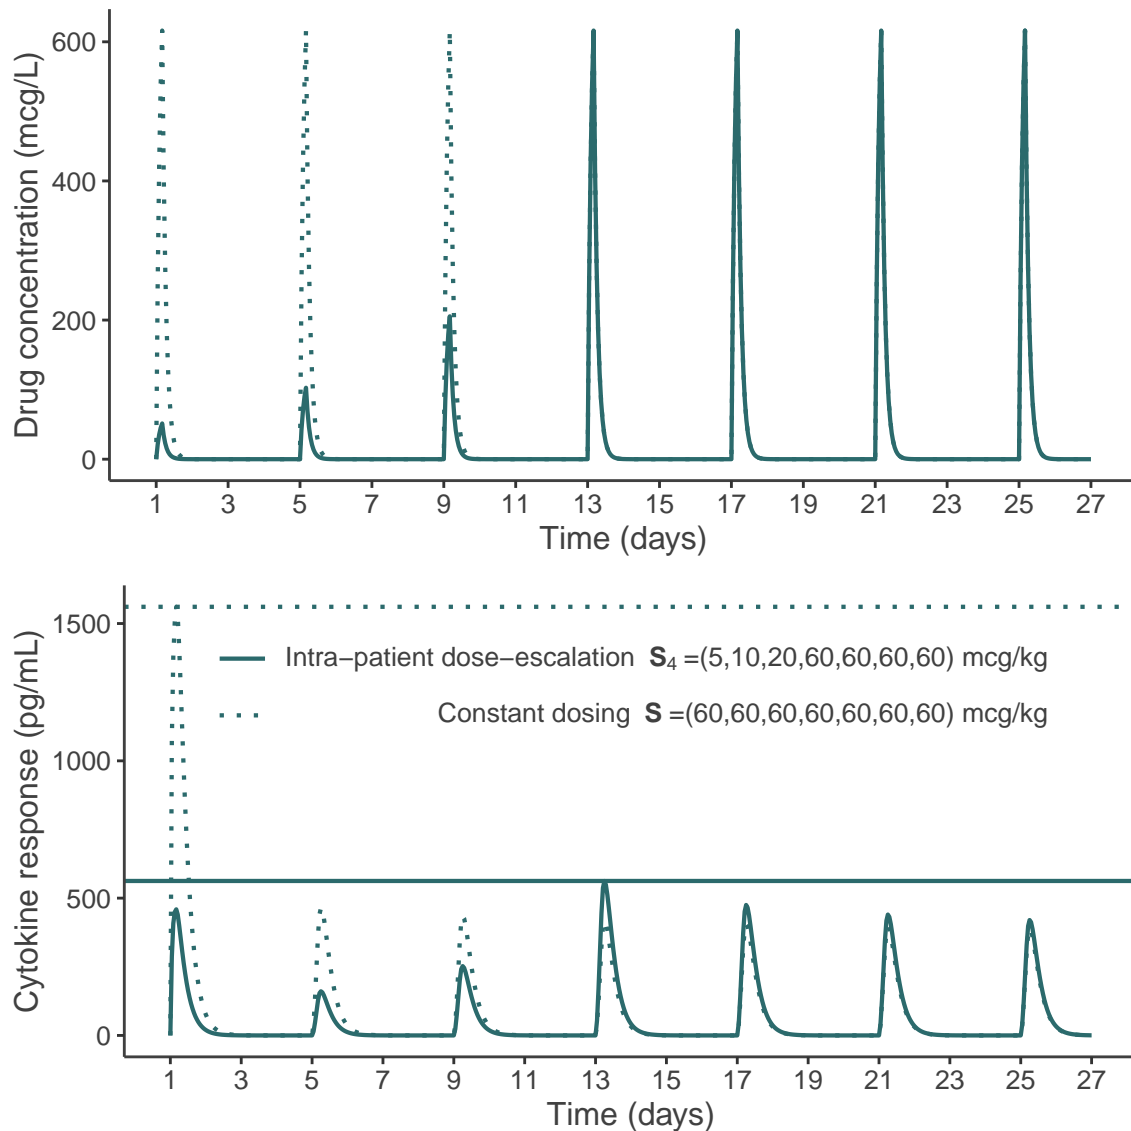


Figure 3.3 – Example of a panel of dose regimens defined with intra-patient dose-escalation (doses in $\mu\text{g}/\text{kg}$).

In Figure 3.3, the population PK/PD profiles after a first dose regimen, defined with intra-patient dose-escalation, and a second dose regimen, when the same dose (the previous steady-state dose) is given repeatedly, are shown. The PK profiles of the two dose regimens are identical from the 4th administration. However, the peak of cytokine of the first dose regimen, defined with constant dosing, that is observed after the first administration is much higher than the peak of cytokine of the second dose regimen, defined with intra-patient dose-escalation.

3.2 Objectives and outline of the manuscript

The objective of this work was to develop a modeling approach in order to estimate the maximum tolerated dose regimen (MTD-regimen), defined as the dose regimen having the DLT probability the closest to the target toxicity probability. We developed this approach in the context of the motivating trial, that is, accounting for the intra-patient dose-escalation scheme implemented to decrease the risk of the main type of DLT expected, CRS, that is assumed to be linked to a PD-endpoint. We therefore proposed to model this PD-related toxicity separately from other DLTs, named DLT_o , as illustrated in Figure 3.4.

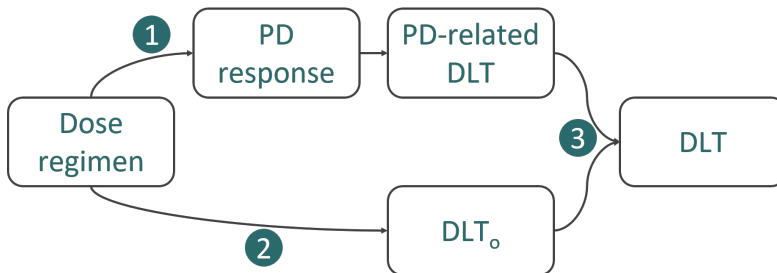


Figure 3.4 – Modeling proposal.

We developed a Bayesian approach, to be applied at the end of the dose-escalation stage of the trial conducted with standard designs, that is divided in three steps:

1. To model the probability of the PD-related toxicity (CRS related to the peak of cytokine in the context of the motivating trial) affected by the intra-patient dose-escalation scheme, we proposed to explicitly model the relationship between the dose regimen and the continuous PD response and the relationship between the PD endpoint and the PD-related toxicity.
2. We then modeled the relationship between the dose regimen and the probability of DLT_o when no prior assumption can be raised on the link with a PD outcome.
3. We finally proposed various approaches to model the joint probability of DLT, defined as a bivariate binary outcome, assuming that the PD-related toxicity and DLT_o can be associated.

All analyses have been performed in R [127] using Stan [16] for Bayesian inference and Monolix [7] for non linear mixed effects models.

The remainder of this thesis manuscript is organized as follows. In Chapter 4, we develop the PD-related toxicity (CRS in the context of the motivating trial) modeling incorporating PK/PD models (step 1 in Figure 3.4), assuming that only this type of DLT can occur, which resulted in a publication in *Biometrics* [64]. In Chapter 5, we develop the modeling of DLT_o with multiple administrations (step 2 in Figure 3.4), in addition to the PD-related outcome, and the different approaches to model the joint distribution of DLT (step 3 in Figure 3.4) which resulted in a publication in *Statistics in Medicine* [?]. Finally, in Chapter 6 we conclude and discuss the work and propose various perspectives.

Chapter 4

Bayesian dose regimen assessment in early phase oncology incorporating pharmacokinetics and pharmacodynamics

In this chapter, we develop the modeling of the relationship between the dose regimen and a PD-related toxicity that resulted in a publication in *Biometrics* [64]. The paper is summarized in Section 4.1 and the complete paper is provided in Section 4.2.

4.1 Summary

Context Standard dose-escalation methods have been developed to find the MTD after the first cycle of therapy accounting for a single administration of treatment. However, patients are usually treated for multiple cycles of therapy and some trials evaluate various schedules of administration in addition to the dose to improve treatment safety while maintaining future potential efficacy.

In the motivating trial, introduced in Section 3.1, intra-patient dose-escalation was implemented to reduce the risk of CRS that is assumed to be linked to a PD endpoint, the peak of cytokine.

The objective of this work was therefore to model the probability of a PD-related toxicity (CRS in the context of the motivating trial) accounting for the complex assumption that states that the multiple administrations that compose the dose regimen decrease the risk of toxicity. Various approaches have been proposed to model repeated administrations in phase I trials as developed in Section 2.2.2 but they do not account for complex toxicity assumptions. One proposal could be to model the probability of each administration with a complex correlation structure between the multiple drug administrations in order to account for this assumption on toxicity. However, this approach can result in a complex modeling approach for early phase trials that have a small sample size. We instead proposed to model the probability of this PD-related toxicity by incorporating PK/PD modeling to account for the complex correlation between multiple administrations. We developed this approach, the DRtox, to be applied at the end of the dose-escalation stage of the trial to determine the MTD-regimen, when all data (especially PK/PD data) was collected.

Method We developed the DRtox approach that describes the relationship between the dose regimen and the probability of the PD-related toxicity by modeling the relationship between the dose regimen and the continuous PD response and the relationship between the PD endpoint and the binary toxicity.

Firstly, the continuous PD response was modeled using nonlinear mixed effects models.

Secondly, the probability of toxicity was modeled with the PD endpoint and we proposed two Bayesian approaches: a Bayesian logistic model (Logistic-DRtox) where the probability of toxicity of the dose regimen was modeled with the global PD endpoint and a Bayesian hierarchical model (Hierarchical-DRtox) where the probability of toxicity after each administration was modeled with the local PD endpoint observed after each administration. We proposed to elicit the prior distributions from the initial guesses of the probabilities of toxicity of each dose regimen. We measured the amount of information provided by the prior distributions using the prior effective sample size (ESS).

Finally, to estimate the probability of toxicity of each dose regimen, we had to integrate the PD endpoint toxicity model on all possible values of the PD endpoint for each dose regimen in order to account for the interindividual variability. We proposed to estimate the posterior probability of toxicity by a Monte Carlo approach and defined the MTD-regimen as the dose regimen having the estimated mean probability of toxicity the closest to the target probability of toxicity.

Results We evaluated the performance of the Logistic-DRtox and Hierarchical-DRtox on an extensive simulation study based on the motivating trial, considering CRS for the PD-related toxicity and the peak of cytokine for the PD endpoint. For the simulations, we used the PK/PD models developed for blinatumomab. As our methods were developed to be applied at the end of the dose-escalation stage of the trial, we compared the results obtained after two standard dose-escalation methods, the 3+3 and a modified CRM. We simulated CRS when the cytokine response of a patient exceeded a threshold taking into account the inter-patient variability.

We observed that our methods had a higher percentage of selecting the correct dose regimen than the standard dose-escalation methods. We also observed that the dose-escalation method implemented had an impact on the results, especially we observed poorest results after a 3+3 design due to the fact that the 3+3 has a tendency to include a small sample size with more patients allocated to sub-optimal dose regimens.

We also evaluated the estimation of the probability of toxicity and observed that, while the CRM and our two proposed methods had a good estimation of the probability of toxicity around the MTD-regimen, both proposed methods had a better estimation of the entire dose-toxicity curve due to the inclusion of PK/PD modeling. Moreover, as our methods modeled the relationship between the entire dose regimen and the probability of toxicity, they could predict the probability of toxicity of untested dose regimens, for example in case the initial panel of dose regimens missed the true MTD-regimen.

Conclusion We proposed an innovative approach to model the relationship between the dose regimen and a PD-related toxicity by incorporating complete PK/PD modeling. However, our modeling proposal has to be applied at the end of the dose-escalation stage of the trial and requires strong assumptions on the PK/PD mechanism generating the toxicity. Moreover, this approach should be performed in close collaboration with pharmacometricians. In this approach, we assumed that the PD-related toxicity was the only type of DLT that could occur, therefore an extension of this modeling proposal was required to account for other toxicities, that are not related to a PD endpoint. This extension is shown in Chapter 5.

4.2 Publication

The supporting information of this paper is provided in Appendix C.

Bayesian dose regimen assessment in early phase oncology incorporating pharmacokinetics and pharmacodynamics

Emma Gerard^{1,2,3}  | Sarah Zohar¹  | Hoai-Thu Thai⁴ | Christelle Lorenzato² | Marie-Karelle Riviere³ | Moreno Ursino^{1,5} 

¹ Centre de Recherche des Cordeliers, Sorbonne Université, Inserm, Université de Paris, Paris F-75006, France

² Oncology Biostatistics, Biostatistics and Programming Department, Sanofi R&D, Vitry-sur-Seine, France

³ Statistical Methodology Group, Biostatistics and Programming Department, Sanofi R&D, Chilly-Mazarin, France

⁴ Translation Disease Modeling, Digital and Data Science, Sanofi R&D, Chilly-Mazarin, France

⁵ F-CRIN PARTNERS Platform, AP-HP, Université de Paris, Paris, France

Correspondence

Sarah Zohar, Centre de Recherche des Cordeliers, Sorbonne Université, Inserm, Université de Paris, F-75006, Paris, France.
Email: sarah.zohar@inserm.fr

Funding information

Association Nationale de la Recherche et de la Technologie, Convention industrielle de formation par la recherche, Grant/Award Number: 2018/0530

Abstract

Phase I dose-finding trials in oncology seek to find the maximum tolerated dose of a drug under a specific schedule. Evaluating drug schedules aims at improving treatment safety while maintaining efficacy. However, while we can reasonably assume that toxicity increases with the dose for cytotoxic drugs, the relationship between toxicity and multiple schedules remains elusive. We proposed a Bayesian dose regimen assessment method (DRtox) using pharmacokinetics/pharmacodynamics (PK/PD) to estimate the maximum tolerated dose regimen (MTD-regimen) at the end of the dose-escalation stage of a trial. We modeled the binary toxicity via a PD endpoint and estimated the dose regimen toxicity relationship through the integration of a dose regimen PD model and a PD toxicity model. For the first model, we considered nonlinear mixed-effects models, and for the second one, we proposed the following two Bayesian approaches: a logistic model and a hierarchical model. In an extensive simulation study, the DRtox outperformed traditional designs in terms of proportion of correctly selecting the MTD-regimen. Moreover, the inclusion of PK/PD information helped provide more precise estimates for the entire dose regimen toxicity curve; therefore the DRtox may recommend alternative untested regimens for expansion cohorts. The DRtox was developed to be applied at the end of the dose-escalation stage of an ongoing trial for patients with relapsed or refractory acute myeloid leukemia (NCT03594955) once all toxicity and PK/PD data are collected.

KEYWORDS

Bayesian inference, dose regimen, early phase oncology, hierarchical model, pharmacokinetics/pharmacodynamics, toxicity

1 | INTRODUCTION

Phase I dose-finding clinical trials in oncology seek to find the maximum tolerated dose (MTD) to obtain reliable information regarding the safety profile of a drug or a combination of drugs, pharmacokinetics, and the mech-

anism of action (Crowley *et al.*, 2005; Chevret, 2006). In this phase, the endpoint is defined as the dose-limiting toxicity, which is mainly based on the National Cancer Institute (NCI) Common Toxicity Criteria for Adverse Events (CTCAE, 2017). Usually, standard algorithm-based or model-based dose-escalation methods (Storer, 1989;

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs License](https://creativecommons.org/licenses/by-nc-nd/4.0/), which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2021 The Authors. *Biometrics* published by Wiley Periodicals LLC on behalf of International Biometric Society.

O'Quigley *et al.*, 1990) aim to find the MTD while considering the entire cycle dosing as a single administration. Most methods assume that toxicity increases with the dose; however, the estimation of the relationship between toxicity and multiple doses over a cycle remains elusive as we can observe nonlinear dose-response profiles (Schmoor and Schumacher, 1992; Bullock *et al.*, 2017; Musuamba *et al.*, 2017). We assume that considering the complete cycle dosage could improve treatment safety while maintaining future potential efficacy.

To account for dosage repetition over the treatment cycle, some authors have considered either the dose-schedule or the dose regimen relationship. The NCI defines "schedule" as "A step-by-step plan of the treatment that a patient is going to receive. A treatment schedule includes the type of treatment that will be given (such as chemotherapy or radiation therapy), how it will be given (such as by mouth or by infusion into a vein), and how often it will be given (such as once a day or once a week). It also includes the amount of time between courses of treatment and the total length of time of treatment" (<https://www.cancer.gov/publications/dictionaries/cancer-terms/def/treatment-schedule>). Moreover, the NCI defines "regimen" as "A treatment plan that specifies the dosage, the schedule, and the duration of treatment" (<https://www.cancer.gov/publications/dictionaries/cancer-terms/def/regimen>). Following these definitions, we considered the dose regimen relationship, as it includes the dosage, the repetition scheme, and the duration.

For some molecules, it has been observed that, in the same patient, starting a dose regimen with a lower lead-in dose and increasing the dose step-by-step before reaching the steady-state dose can reduce the occurrence of acute toxicities (Chen *et al.*, 2019). However, a dose regimen starting with higher lead-in doses can increase the efficacy.

Dose-finding trials can aim to study different dose regimens with the same or different total cumulative dose to determinate the most appropriate regimen supported by pharmacokinetics/pharmacodynamics (PK/PD) profiles. Several methodological papers have attempted to address the issue of prospective dose and schedule finding methods. Braun *et al.* (2005), Braun *et al.* (2007), Liu and Braun (2009), and Zhang and Braun (2013) proposed considering the time-to-toxicity rather than the usual binary outcome to optimize dose and schedule, as the timing of administration. Wages *et al.* (2014) proposed considering dose-schedule finding as a two-dimensional problem and extended the partial-order continual reassessment method developed for combination trials. Other authors, such as Li *et al.* (2008), Thall *et al.* (2013), and Guo *et al.* (2016), proposed dose-schedule-finding methods that jointly model toxicity and efficacy outcomes. Lyu *et al.* (2018) proposed a hybrid design that is partially algorithm-based and par-

tially model-based for sequences of doses over multiple cycles when few doses are under study.


Only a few methods consider PK/PD data in the prospective dose-allocation design. Ursino *et al.* (2017) compared multiple methods that enable the use of PK measures in sequential Bayesian adaptive dose-finding designs, including a dose-AUC-toxicity model combining two models to recommend the dose. Günhan *et al.* (2020) proposed a Bayesian time-to-event pharmacokinetic adaptive model for multiple regimens using PK latent profiles to measure drug exposure. Our aim is to extend these propositions by modeling the dose regimen toxicity relationship using PK/PD.

2 | MOTIVATION

This work was motivated by the ongoing first-in-human dose-escalation study of SAR440234 (<https://www.cancer.gov/publications/dictionaries/cancer-drug/def/798327>) administered as a single agent to patients with relapsed or refractory acute myeloid leukemia, high-risk myelodysplastic syndrome, or B-cell acute lymphoblastic leukemia (NCT03594955 <https://clinicaltrials.gov/ct2/show/NCT03594955>). SAR440234 is a novel bispecific T-cell engager antibody that activates and redirects cytotoxic T lymphocytes (CTLs) to enhance the CTL-mediated elimination of CD123-expressing tumor cells. CTL activation induces the release of inflammatory cytokines, which can potentially cause cytokine release syndrome (CRS). CRS is a systemic inflammatory response and among the most commonly observed toxicities of T-cell engaging bispecific antibodies, such as blinatumomab, which is a bispecific anti-CD19/CD3 antibody (Shimabukuro-Vornhagen *et al.*, 2018). Several cytokines, such as IL6, IL10, and $\text{INF}\gamma$, are consistently found to be elevated in serum from patients with CRS. The association between the peak of cytokine and CRS has been evaluated by Teachey *et al.* (2016). It has been shown that repeating the dosing of the drug can decrease CRS, particularly when the first administration is divided into several steps progressively (Chen *et al.*, 2019). Therefore, inpatient dose-escalation with a dose regimen consisting of lower initial doses followed by a higher maintenance dose was implemented in this study to reduce the occurrence of CRS (Boissel *et al.*, 2018).

The aim of the trial was to find the MTD of SAR440234 using the 3+3 design as the dose-escalation design. However, the 3+3 design and more general dose-finding designs ignore part of the dose regimen information: these designs were not developed to account for multiple dose administrations in the model. Therefore, they map the entire dose regimen administered to the patient in a single dose-level, that is, a single value. This mapping is defined prior to

FIGURE 1 Trial scheme: the DRtox method is applied at the end of the dose-escalation stage of a phase I trial



the trial onset and depends on the design chosen. This approach is inefficient for achieving the trial goal.

In this paper, we propose to model the binary toxicity endpoint (CRS) and the continuous PD response (cytokine profile) at the end of the trial, once all data have been collected, to characterize the dose regimen toxicity relationship. This dose regimen assessment method (DRtox) allows the determination of the maximum tolerated dose regimen (MTD-regimen), as illustrated in Figure 1.

3 | MODEL

Let $\mathcal{D} = \{d_1, \dots, d_L\}$ be the set of doses that can be administered to patients, where $d_l < d_{l+1}$. Let $\mathcal{S} = \{\mathcal{S}_1, \dots, \mathcal{S}_K\} \subset \mathcal{S}$ be the panel of dose regimens to be studied in the trial. The dose regimen $\mathcal{S}_k \in \mathcal{S}$, where $k \in \{1, \dots, K\}$, is defined as the sequence of J doses, $\mathcal{S}_k = (d_{k,1}, d_{k,2}, \dots, d_{k,J})$, administered at times $\mathbf{t} = (t_1, t_2, \dots, t_J)$, where $d_{k,j} \in \mathcal{D}$ for $j \in \{1, \dots, J\}$. To simplify the notations, we assumed that all dose regimens have the same number of drug administrations at the same times, but this assumption can be relaxed. Let $\mathcal{S}_{k,j}$ be the subregimen of \mathcal{S}_k until the j th administration, $\mathcal{S}_{k,j} = (d_{k,1}, d_{k,2}, \dots, d_{k,j})$, for $j < J$. Let $n \in \mathbb{N}$ be the number of patients included in the trial. Let $Y_{i,j}$ be the binary toxicity response of patient i observed exactly after the j th administration, and let Y_i be his/her global toxicity response at the end of the administrations.

Let $\tilde{\mathbf{s}}_i = (d_{i,1}, d_{i,2}, \dots, d_{i,J}) \in \mathcal{S}$ be the dose regimen planned for the i th patient. We assume that the drug administration is stopped if toxicity occurs; thus let j_i denote the last administration of patient i . We denote the actual regimen received by patient i as $\mathbf{s}_i = (d_{i,1}, d_{i,2}, \dots, d_{i,j_i}) \subset \tilde{\mathbf{s}}_i$, where $\mathbf{s}_i = \tilde{\mathbf{s}}_i$ if no toxicity is observed. Let $\mathbf{s}_{i,j}$ be the subregimen until j of \mathbf{s}_i , where $j \leq j_i$.

The aim is to estimate the MTD-regimen at the end of the trial, which is defined as the dose regimen with the toxicity probability closest to the target toxicity rate δ_T , that is, the MTD-regimen is the regimen \mathcal{S}_{k^*} , where $k^* = \arg \min_k |p_T(\mathcal{S}_k) - \delta_T|$ and $p_T(\mathcal{S}_k)$ is the toxicity probability of \mathcal{S}_k .

We assume that a PD endpoint extracted from the continuous PD profile of a biomarker related to toxicity plays an intermediate role in the dose regimen toxicity relationship. We propose the DRtox approach in which the first

model is built for the dose regimen and the PD endpoint, and the second model is built for the PD endpoint and the toxicity response. Therefore, integrating both models links the dose regimen to the toxicity response to find the MTD-regimen. In the following section, the structure of the PK/PD models is described, two approaches between the PD endpoint and toxicity response are proposed, as well as a practical method for their integration.

3.1 | Dose regimen PD response model

Let $C(t)$ be the continuous drug concentration and $E(t)$ be the continuous PD response related to toxicity measured at time t . We assume that $C(t)$ and $E(t)$ can be modeled using nonlinear mixed-effects models as follows:

$$\begin{cases} C(t) = f^{(1)}(\theta_i^{(1)}, t) + g^{(1)}(\theta_i^{(1)}, t, \xi_1) \epsilon^{(1)}, \\ E(t) = f^{(2)}(\theta_i^{(2)}, t) + g^{(2)}(\theta_i^{(2)}, t, \xi_2) \epsilon^{(2)}, \end{cases} \quad (1)$$

where $f^{(1)}$ and $f^{(2)}$ represent the structural models, which are usually solutions of differential equations based on biological knowledge. $\theta_i = (\theta_i^{(1)}, \theta_i^{(2)})$ represents the i th patient's specific parameter vector, where usually, $\theta_i = \mu e^{\eta_i}$, with μ denoting the fixed effects vector, and η_i denoting the random effects vector defined as $\eta_i \sim \mathcal{N}(\mathbf{0}, \mathbf{\Omega})$, with $\mathbf{\Omega}$ denoting the variance-covariance matrix.

$g^{(1)}$ and $g^{(2)}$ represent the error models, which depend on the additional parameters ξ_1 and ξ_2 , and $\epsilon^{(1)}$ and $\epsilon^{(2)}$ are standard Gaussian variables. The usual error models are the constant model where $g^{(l)}(\theta_i^{(l)}, t, \xi_l = a) = a$, the proportional model where $g^{(l)}(\theta_i^{(l)}, t, \xi_l = b) = b f^{(l)}(\theta_i^{(l)}, t)$ and combinations of the constant and proportional models.

3.2 | PD endpoint toxicity model

$r(\theta_i, \mathbf{s}_{i,j})$ is defined as the function derived from the PK/PD models that returns the value of the PD endpoint (such as the peak of a biomarker) exactly after the administration of the dose regimen $\mathbf{s}_{i,j}$ with individual PK/PD parameters θ_i . Let $\mathbf{R}(\theta_i, \mathbf{s}_{i,j}) = (r(\theta_i, \mathbf{s}_{i,1}), \dots, r(\theta_i, \mathbf{s}_{i,j}))$ be the function derived from the PK/PD models that returns the vector of all PD endpoints (such as all biomarker peaks) observed

after the administration of the regimen $\mathbf{s}_{i,j}$ with individual PK/PD parameters θ_i . For patient i , we can simplify the notations considering $r_{i,j} = r(\theta_i, \mathbf{s}_{i,j})$, $\mathbf{R}_{i,j} = \mathbf{R}(\theta_i, \mathbf{s}_{i,j})$ and the vector of all PD endpoints $\mathbf{R}_i = \mathbf{R}_{i,j_i}$.

Then, let $r_i^M = \max_{l \in \{1, \dots, j_i\}} (r_{i,l})$ be the summary PD endpoint (such as the highest peak) observed in patient i , which we assume is related to toxicity.

To define the prior distributions, let $(\bar{r}_1^M, \bar{r}_2^M, \dots, \bar{r}_K^M)$ denote the reference values of the summary endpoint of all dose regimens of the trial $(\mathbf{S}_1, \dots, \mathbf{S}_K)$; for example, we can consider population values $\bar{r}_k^M = \max\{r(\boldsymbol{\mu}, \mathbf{S}_{k,1}), \dots, r(\boldsymbol{\mu}, \mathbf{S}_k)\}$ with $\boldsymbol{\mu}$ as the PK/PD vector of fixed effects.

In the following section, two statistical models establishing the relationship between the PD endpoint and the toxicity response are shown.

3.2.1 | Logistic-DRtox

We propose a Bayesian logistic model to link the global binary toxicity response of patient i receiving \mathbf{s}_i to his summary PD endpoint related to toxicity as follows:

$$\text{logit}\{\mathbb{P}(Y_i = 1)\} = \beta_0 + \beta_1 \log\left(\frac{r_i^M}{\bar{r}_{k_T}^M}\right), \quad (2)$$

where $\beta_1 > 0$ to have the toxicity probability that increases with the value of the summary PD endpoint. We normalize the PD endpoint for prior elicitation using $\bar{r}_{k_T}^M$, which is the reference value of dose regimen \mathbf{S}_{k_T} , which we initially guess to have a toxicity probability of δ_T . In this model, we do not consider the longitudinal values of the biomarker as we assume that toxicity is not due to the cumulative effect of the biomarker profile. However, previous drug administrations are considered in the construction of the biomarker through the PK/PD model. Let $\pi_1\{(\beta_0, \beta_1), r_i^M\} = \text{logit}^{-1}\left\{\beta_0 + \beta_1 \log\left(\frac{r_i^M}{\bar{r}_{k_T}^M}\right)\right\}$.

Regarding prior distributions, we consider a normal distribution for the intercept, $\beta_0 \sim \mathcal{N}(\bar{\beta}_0, \sigma_{\beta_0}^2)$ and a gamma distribution for the slope to ensure positivity, $\beta_1 \sim \gamma(\alpha_1, \frac{\alpha_1}{\bar{\beta}_1})$, where α_1 is the shape parameter, $\bar{\beta}_0 = \mathbb{E}[\beta_0]$, and $\bar{\beta}_1 = \mathbb{E}[\beta_1]$. By construction, we have $\bar{\beta}_0 = \text{logit}(\delta_T)$, obtained via Equation (2) with $r_i^M = \bar{r}_{k_T}^M$. Then, let (p_1, \dots, p_K) be the initial guesses of the toxicity probabilities of regimens $(\mathbf{S}_1, \dots, \mathbf{S}_K)$, where $p_{k_T} = \delta_T$. We can determine $\bar{\beta}_1$ using either only one regimen, which differs from the reference regimen \mathbf{S}_{k_T} , as $\pi_1\{(\bar{\beta}_0, \bar{\beta}_1), \bar{r}_k^M\} = p_k$, with $k \in \{1, \dots, K\}$ and $k \neq k_T$, or multiple regimens, such

as the neighbors of the reference regimen, as follows:

$$\bar{\beta}_1 = \arg \min_{\beta_1} \sum_{k=k_T-1}^{k_T+1} \left[p_k - \pi_1\left\{(\bar{\beta}_0, \beta_1), \bar{r}_k^M\right\} \right]^2. \quad (3)$$

3.2.2 | Hierarchical-DRtox

In this approach, we assume that patients experience toxicity if their PD response exceeds an unknown threshold specific to each patient. To consider interindividual variability in toxicity, we introduce a patient-specific continuous latent variable, Z_i , which represents the toxicity threshold of the PD response. In contrast to the previous approach, we model toxicity after each administration using a modification of the hierarchical probit model (Berry *et al.*, 2010) as follows:

$$\begin{cases} Y_{i,j} = \begin{cases} 0 & \text{if } Z_i > \log\left(\frac{r_{i,j}}{\bar{r}_{k_{50}}^M}\right) \\ 1 & \text{if } Z_i \leq \log\left(\frac{r_{i,j}}{\bar{r}_{k_{50}}^M}\right) \end{cases} \\ Z_i \sim \mathcal{N}(\mu_z, \tau_z^2), \end{cases} \quad (4)$$

where $\bar{r}_{k_{50}}^M$ is the reference value at the dose regimen $\mathbf{S}_{k_{50}}$, which we initially guess to have a toxicity probability of 0.5. By adding the random effect, this Bayesian hierarchical model shares common features with the probit model, where τ_z^2 represents the between-subject variance and controls the extent of the borrowing across all patients.

If we consider a new patient i with a vector of biomarker endpoints \mathbf{R}_i , we can predict his probability of toxicity by $\mathbb{P}(Y_i = 1) = F_z\left\{\log\left(\frac{r_i^M}{\bar{r}_{k_{50}}^M}\right)\right\}$, where F_z is the cumulative distribution function of $\mathcal{N}(\mu_z, \tau_z^2)$. The details of the formula are shown in Web Appendix A. Let $\pi_2\{(\mu_z, \tau_z^2), r_i^M\} = F_z\left\{\log\left(\frac{r_i^M}{\bar{r}_{k_{50}}^M}\right)\right\}$.

Regarding the prior distributions, we consider $\mu_z \sim \mathcal{N}(0, \sigma_{\mu_z}^2)$ and $\tau_z \sim \text{half-Cauchy}(0, \sigma_{\tau_z}^2)$. Regarding the half-Cauchy distribution, we followed the recommendations by Gelman (2006), as we assumed that τ_z could be near 0. Web Appendix G shows how this model can be implemented.

3.3 | Dose regimen toxicity model

The posterior toxicity probability of dose regimen \mathbf{S}_k is estimated by integrating the PD endpoint toxicity model

on all possible values of the PD endpoint. As this integral cannot usually be solved analytically, the posterior toxicity probability of regimen \mathbf{S}_k is estimated via the drawing of a hypothetical set of M patients with M -vector $(p_T(\mathbf{S}_k)^{(1)}, \dots, p_T(\mathbf{S}_k)^{(M)})$ as posterior toxicity probabilities. Then, the posterior toxicity probability of regimen \mathbf{S}_k is estimated as the posterior mean $\widehat{p}_T(\mathbf{S}_k) = \frac{1}{M} \sum_{m=1}^M p_T(\mathbf{S}_k)^{(m)}$. This sample of the posterior toxicity probability requires the following three major steps:

(1) *Model fitting:*

- (a) First, the PK/PD models are fitted to obtain estimates of the population parameters comprising the fixed effects, $\widehat{\boldsymbol{\mu}}$, and the random effects variance–covariance matrix, $\widehat{\boldsymbol{\Omega}}$, under the Frequentist paradigm. The patients' individual PK/PD parameters, $(\widehat{\boldsymbol{\theta}}_1, \dots, \widehat{\boldsymbol{\theta}}_n)$, are also estimated.
- (b) Based on the estimated PK/PD parameters, the PD biomarkers are predicted for each patient:
 - For the logistic-DRtox: the global biomarker peaks $(\widehat{r}_1^M, \dots, \widehat{r}_n^M)$ are predicted for each patient as $\widehat{r}_i^M = \max\{r(\widehat{\boldsymbol{\theta}}_i, \mathbf{s}_{i,1}), \dots, r(\widehat{\boldsymbol{\theta}}_i, \mathbf{s}_i)\}$ for $i \in \{1, \dots, n\}$. The vector of toxicity responses and biomarker responses, $((Y_1, \dots, Y_n), (\widehat{r}_1^M, \dots, \widehat{r}_n^M))$, constitutes the data of the trial.
 - For the hierarchical-DRtox: the biomarker peaks vectors $(\widehat{\mathbf{R}}_1, \dots, \widehat{\mathbf{R}}_n)$ are predicted for each patient as $\widehat{\mathbf{R}}_i = \mathbf{R}(\widehat{\boldsymbol{\theta}}_i, \mathbf{s}_i)$ for $i \in \{1, \dots, n\}$. The vector of toxicity responses and biomarker responses, $((Y_{1,1}, \dots, Y_{n,j_n}), (\widehat{\mathbf{R}}_1, \dots, \widehat{\mathbf{R}}_n))$, constitutes the data of the trial.
- (c) A vector of the parameters of the PD endpoint toxicity model of size m_{iter} is sampled from their posterior distribution:
 - For the logistic-DRtox, $((\beta_0^{(1)}, \beta_1^{(1)}), \dots, (\beta_0^{(m_{\text{iter}})}, \beta_1^{(m_{\text{iter}})}))$ is sampled.
 - For the hierarchical-DRtox, $((\mu_z^{(1)}, \tau_z^{(1)}), \dots, (\mu_z^{(m_{\text{iter}})}, \tau_z^{(m_{\text{iter}})}))$ is sampled.

(2) *Prediction of new patients for the sampling distribution of the PD endpoint:*

- (a) The individual PK/PD parameters of m_{predict} simulated patients, $(\boldsymbol{\theta}^{(1)}, \dots, \boldsymbol{\theta}^{(m_{\text{predict}})})$, are sampled from $\widehat{\boldsymbol{\mu}}$ and $\widehat{\boldsymbol{\Omega}}$ as $\boldsymbol{\theta}^{(m_p)} = \widehat{\boldsymbol{\mu}} e^{\boldsymbol{\eta}^{(m_p)}}$, with $\boldsymbol{\eta}^{(m_p)} \sim \mathcal{N}(\mathbf{0}, \widehat{\boldsymbol{\Omega}})$ for $m_p \in \{1, \dots, m_{\text{predict}}\}$.
- (b) The maximum biomarker endpoint of each simulated patient receiving regimen \mathbf{S}_k is obtained as $r^{(m_p)} = \max(r(\boldsymbol{\theta}^{(m_p)}, \mathbf{S}_{k,1}), \dots, r(\boldsymbol{\theta}^{(m_p)}, \mathbf{S}_k))$ for $m_p \in \{1, \dots, m_{\text{predict}}\}$.

(3) *Estimation of the posterior distribution of the probability of toxicity:*

- (a) The m th iteration, $m = (m_i, m_p) \in \{1, \dots, M\}$, where $M = m_{\text{iter}} * m_{\text{predict}}$, of the posterior probability of toxicity of dose regimen \mathbf{S}_k , $p_T(\mathbf{S}_k)^{(m)}$, is obtained depending on the method chosen:
 - For the logistic-DRtox, $p_T(\mathbf{S}_k)^{(m)} = \pi_1 \left\{ \left(\beta_0^{(m_i)}, \beta_1^{(m_i)} \right), r^{(m_p)} \right\}$.
 - For the hierarchical-DRtox, $p_T(\mathbf{S}_k)^{(m)} = \pi_2 \left\{ \left(\mu_z^{(m_i)}, \tau_z^{(m_i)} \right), r^{(m_p)} \right\}$.

The DRtox approach allows us to estimate the toxicity probability of the panel of dose regimens \mathbf{S} and predict the toxicity probability of each new regimen defined from the set of doses \mathbf{D} .

4 | SIMULATION STUDY

4.1 | Simulation settings

The performance of the DRtox was evaluated through a simulation study. We assumed that toxicity was related to a PD endpoint (the peak of cytokine in the context of our motivating example). Therefore, to simulate toxicity, we first needed to simulate the PK/PD profiles and simulate toxicity from the PD profile.

Regarding the PK/PD models, we were inspired by published models on blinatumomab, which is another bispecific T-cell engager that binds to CD3 on T-cells and to CD19 on tumor cells. Regarding the PK model, we considered a 1-compartment infusion model (Zhu *et al.*, 2016) in which the parameters are the volume of distribution V and the clearance of elimination Cl and assumed 4 h of infusion. The model is defined in Web Appendix B. Regarding the PD aspect, the objective was to model cytokine mitigation in the case of inpatient dose-escalation. We simplified the model developed by Chen *et al.* (2019), which assumes that cytokine production is stimulated by the drug concentration but inhibited by cytokine exposure through the AUC. We defined the PD model as follows:

$$\frac{dE(t)}{dt} = \frac{E_{\max} C(t)^H}{EC_{50}^H + C(t)^H} \left[1 - \frac{I_{\max} AUC_E(t)}{\frac{IC_{50}}{K^{J-1}} + AUC_E(t)} \right] - k_{\text{deg}} E(t), \quad (5)$$

where $E(t)$ and $C(t)$ are the cytokine and drug concentration at time t , respectively, $AUC_E(t)$ is the cumulative cytokine exposure, and the parameters are defined in Table 1. Additional information concerning the PK/PD models is provided in Web Appendix B.

In both the PK and PD models, we considered a proportional error model with $b = 0.1$. The values of the PK/PD

TABLE 1 Definition and values of the PK/PD parameters used for the simulation study. Parameter estimates represent the fixed effects, and coefficients of variation (CV) are the square root of the diagonal of the variance–covariance matrix. They are inspired by the parameters estimated on blinatumomab (Zhu *et al.*, 2016; Chen *et al.*, 2019), with a modification of I_{\max} to observe cytokine mitigation after several administrations

	Parameter	Estimate (% CV)	Unit	Description
PK model	Cl	1.36 (41.9)	L/h	Clearance
	V	3.4 (0)	L	Volume of distribution
PD model	E_{\max}	$3.59 \cdot 10^5$ (14)	pg/mL/h	Maximum cytokine release rate
	EC_{50}	$1 \cdot 10^4$ (0)	ng/mL	Drug exposure for half-maximum cytokine release rate
	H	0.92 (3)		Hill coefficient for cytokine release
	I_{\max}	0.995 (0)		Maximum inhibition of cytokine release
	IC_{50}	$1.82 \cdot 10^4$ (12)	pg/mL·h	Cytokine exposure for half-maximum cytokine inhibition
	k_{deg}	0.18 (13)	h^{-1}	Degradation rate for cytokine
	K	2.83 (36)		Priming factor for cytokine release

parameters used for the simulations were inspired by the estimated parameters of blinatumomab (Zhu *et al.*, 2016; Chen *et al.*, 2019) and are displayed in Table 1.

To simplify and accelerate the PK/PD estimation during the simulations, we followed the traditional PK/PD modeling strategy for small sample size data by fixing some parameters. We considered the parameters EC_{50} , I_{\max} , and IC_{50} fixed and no random effects on V and H. In this work, we decided to simplify a previously validated PK/PD model that mimics the behavior we expect in our motivating trial: our aim was to show the performance of a global modeling approach including PK/PD estimation in a phase I toxicity model and not to propose a PK/PD model for the drug.

We used as the PD endpoint $r_{i,j}$ the peak of cytokine observed for patient i after the j th administration, and for r_i^M the highest peak of cytokine observed for patient i . Using the PK/PD models presented above and the parameters shown in Table 1, we were able to model the mitigation of cytokine release on repeating dosing, which was reflected by the decrease in the cytokine peak with repeating dosing. Hence, we were able to model that slowly increasing the dose reduces the cytokine peak compared to directly giving the steady-state dose. For example, we compared the concentration and cytokine profiles of patients i and i' who received regimens $\mathbf{s}_i = (1, 5, 10, 25, 25, 25, 25)$ $\mu\text{g}/\text{kg}$ and $\mathbf{s}_{i'} = (25, 25, 25, 25, 25, 25, 25)$ $\mu\text{g}/\text{kg}$ administered on days 1, 5, 9, 13, 17, 21, and 25 (Figure 2). From the fourth administration, the concentration profiles of patients i and i' are the same, but in the cytokine profile, the maximum peak of cytokine of patient i' is much higher than that of patient i , $r_{i'}^M = r_{i',1}^M > r_i^M = r_{i,4}^M$.

To simulate toxicity from the cytokine profile, we defined a threshold τ_T on the cytokine response and assumed that toxicity occurred if this threshold was exceeded (Ursino *et al.*, 2017). To introduce between-subject variability, we defined a log-normally distributed

measure of subject sensitivity, α_i for patient i , where $\alpha_i = e^{\eta_{\alpha_i}}$ and $\eta_{\alpha_i} \sim \mathcal{N}(0, \omega_{\alpha}^2)$. We assumed that patient i experienced toxicity at the j th administration, $Y_{i,j} = 1$, if $\alpha_i r_{i,j} \geq \tau_T$.

To compute the toxicity probability of regimen \mathbf{S}_k , we used the Monte-Carlo method by simulating $N = 10,000$ cytokine profiles under \mathbf{S}_k and computing

$$p_T(\mathbf{S}_k) = \frac{1}{N} \sum_{i=1}^N \left[1 - \Phi \left\{ \frac{\log(\tau_T) - \log(r_i^M)}{\omega_{\alpha}} \right\} \right], \quad (6)$$

where Φ is the cumulative distribution function of the standard normal distribution.

We present the results of three toxicity scenarios by varying the dose regimens and the value of the threshold τ_T to explore different positions of the MTD-regimen (with $\omega_{\alpha} = 0.25$). Additional scenarios are presented in Web Appendix F. In each scenario, we considered six dose regimens, and each dose regimen included seven dose administrations on days 1, 5, 9, 13, 17, 21, and 25. The dose regimens chosen for each scenario and the dose regimen toxicity curves are displayed in Figure 3. Values of the dose regimens can be found in Web Appendix C. In Scenarios 1–3, the MTD-regimen is situated at dose regimens \mathbf{S}_4 , \mathbf{S}_2 , and \mathbf{S}_4 , respectively. Scenarios 1 and 2 are inspired from the motivating trial in which the dose regimens reach the steady-state dose at approximately the same time, and have increasing steady-state doses. However, Scenario 3 represents a case in which the objective is to reach the steady-state dose, 40 $\mu\text{g}/\text{kg}$, as fast as possible to increase potential efficacy under toxicity constraints. The dose regimen toxicity relationship is similar to that in Scenario 1 but with less difference between the MTD-regimen and its neighbors.

For each scenario, 1000 trials were simulated, and $\delta_T = 0.3$ was considered the toxicity target. Because we applied our methods once all patients from the trial were included,




FIGURE 2 Concentration (up) and cytokine (down) profiles of two patients, one receiving a dose regimen with inpatient escalation in solid line and the other receiving a dose regimen without inpatient escalation in dashed line, administered on days 1, 5, 9, 13, 17, 21, and 25. Horizontal lines represent the maximum peak of cytokine observed after each dose regimen

we evaluated the impact of two traditional dose-escalation designs, that is, the 3+3 design and a modified continual reassessment method (CRM) initially proposed by O’Quigley *et al.* (1990). A flow diagram of the rules of the 3+3 design is provided in Web Appendix D. For the modified CRM, we considered a 2-parameter logistic regression model with cohorts of a size of 3 and a total sample size of 30 patients (Cheung, 2011). Dose skipping was not allowed, and early stopping rules were not implemented. We based the skeleton of the CRM, that is, the prior guesses of the toxicity probabilities, on Scenario 1, that is, (0.06, 0.12, 0.20, 0.30, 0.40, 0.50). This skeleton was used in all simulations and scenarios.

When defining the prior distributions for our proposed models, we calibrated the model prior distributions based on the initial guesses of the toxicity probabilities (we used the same initial guesses for the CRM). To quantify the information provided by the prior distribution, we computed the approximate effective sample size (ESS), which was defined as the equivalent sample size embedded in the prior distribution of the model parameters (Yuan *et al.*, 2017). In practice, we approximated the ESS by matching the mean and variance of the toxicity probabilities computed from the prior distributions to those of a beta distribution. Then, the ESS was computed as the sum of the parameters of the beta distribution (Morita *et al.*, 2008). More details of the ESS computation are shown in Web Appendix E. In our settings, for the logistic-DRtox, we considered $k_T = 4$, $\sigma_{\beta_0} = 2$, and $\alpha = 5$, leading to an approxi-

mate mean ESS of 1.6. For the hierarchical DRtox, we considered $k_{50} = 6$, $\sigma_{\mu} = 1$, and $\sigma_{\tau} = 1$, leading to an approximate mean ESS of 1.8.

All simulations were performed in the R environment (R Core Team, 2018), using Monolix software (Lixoft SAS, 2019) for the PK/PD estimation and Stan (Stan Development Team, 2019) for the Bayesian analysis.

4.2 | Simulation results

4.2.1 | Proportions of correct selection

We first evaluated the performance of the DRtox according to the proportions of correct selection (PCS) based on the proportions that each regimen is selected as the MTD-regimen over the trials. We evaluated the impact of the dose regimens and the position of the MTD-regimen in three toxicity scenarios, and the impact of the dose-escalation design, that is, either the 3+3 design or the CRM. The PCS results of the three main scenarios and the mean sample size of each dose regimen across the trials due to the chosen dose-escalation design are displayed in Table 2. The PCS of additional scenarios are displayed in Web Appendix F. As a practical rule, we could only recommend as the MTD-regimen a dose regimen that was administered during the dose-escalation phase of the trial.

In all scenarios, the PCS of the logistic-DRtox and the hierarchical-DRtox are very similar. Both methods




FIGURE 3 The first three subplots represent the panel of dose regimens from S_1 in spaced dashed line to S_6 in solid line, for the three main scenarios, where the type of points is specific to each scenario. In the last subplot in the lower right corner, the dose regimen toxicity relationship is represented for each scenario, where the MTD-regimen is the dose regimen having the toxicity probability the closest to the target δ_T , plotted in dashed line

outperform the dose-escalation design implemented in most scenarios. After implementing the 3+3 design, our methods correctly select the MTD-regimen in more than 10% more trials compared to the dose-allocation design. After implementing the CRM design, both methods correctly select the MTD-regimen in approximately 10% more trials compared to the CRM.

The results of Scenarios 1, 3, and 6 (presented in Web Appendix F) illustrate the effect of the variation in the dose regimen scheme with a similar dose regimen toxicity relationship. Compared to Scenario 1, the PCS of the logistic-DRtox and hierarchical-DRtox are decreased by approximately 10% in Scenario 3, while there is not much difference in the results between Scenarios 1 and 6. Therefore, the loss of performance in Scenario 3 is caused not only by the variation in the dose regimen scheme but also by the difference in the dose regimen toxicity relationship, as in Scenario 3 there is less difference in the toxicity probabilities between the MTD-regimen and its neighbors.

However, the performance of the DRtox is heavily impacted by the dose-escalation design implemented; after

implementing the CRM design, the DRtox correctly selects the MTD-regimen in more than 50% of the trials, but its PCS can decrease by 20% when applied after the 3+3 design. This loss of performance is due to the small sample size after implementing the 3+3 design and the higher proportion of patients allocated to suboptimal dose regimens.

Additional results on robustness (with various prior distributions and variability to simulate toxicity) are given in Web Appendix F.

4.2.2 | Estimation of the toxicity probabilities

We also evaluated the performance of the DRtox based on the precision of the estimation of the toxicity probabilities of all dose regimens. We represented the distribution of the estimated toxicity probabilities, defined as the mean of the posterior distribution, over 1000 trials. The results of Scenario 3 obtained after implementing the CRM are presented in the lower part of Figure 4. The results of the other scenarios are displayed in Web Appendix F.

TABLE 2 Proportions that each dose regimen is being selected as the MTD-regimen over the 1000 trials in the three toxicity scenarios and the two dose-allocation designs, either the 3+3 design or the CRM. For each scenario, the PCS on the true MTD-regimen are represented in bold. For each dose-allocation design, the mean sample size of each dose regimen is displayed

		S_1	S_2	S_3	S_4	S_5	S_6
Scenario 1		0.08	0.11	0.15	0.3	0.44	0.52
3+3	Mean sample size	3.6	3.5	3.5	3	1.6	0.4
	Logistic-DRtox	8.6	5.9	19	42.2	19.6	4.7
	Hierarchical-DRtox	7.5	7.6	19.1	43.8	18.6	3.4
	3+3	13.9	16.1	32.2	27.6	8.6	1.6
CRM	Mean sample size	4.2	3.7	5.6	8.8	5.6	2.1
	Logistic-DRtox	0	1.2	15.5	64.6	15.5	3.2
	Hierarchical-DRtox	0	0.8	12.8	64.3	19.4	2.7
	Logistic CRM	0	1.4	15.1	50.4	27.1	6
Scenario 2		0.15	0.3	0.44	0.52	0.69	0.83
3+3	Mean sample size	4	3.6	1.8	0.5	0.1	0
	Logistic-DRtox	27.2	42.5	24.7	5.2	0.4	0
	Hierarchical-DRtox	29.3	41.2	24.3	4.8	0.4	0
	3+3	57.3	31	9.8	1.7	0.2	0
CRM	Mean sample size	8.7	11.1	7.5	2.3	0.3	0
	Logistic-DRtox	14.8	65.9	17.4	1.7	0.2	0
	Hierarchical-DRtox	12.3	66.2	18.9	2.6	0	0
	Logistic CRM	12.5	56	26.7	4.7	0.1	0
Scenario 3		0.07	0.11	0.2	0.3	0.42	0.56
3+3	Mean sample size	3.6	3.6	3.7	2.7	1.4	0.4
	Logistic-DRtox	7.8	6.4	25.2	34.1	21.6	4.9
	Hierarchical-DRtox	5.9	7.9	27.3	35.8	20.6	2.5
	3+3	13.1	24.4	29.5	24	7.7	1.3
CRM	Mean sample size	4	4	6.4	8	5.2	2.3
	Logistic-DRtox	0.1	1.4	19.6	52	25.1	1.8
	Hierarchical-DRtox	0.1	0.8	17.7	54.4	25.9	1.1
	Logistic CRM	0.1	2.3	20.3	44.5	26.4	6.4

In all scenarios, the toxicity probability of the MTD-regimen is well estimated by the DRtox and the CRM. Both the hierarchical-DRtox and the logistic-DRtox seem to be better in estimating the toxicity probability at all dose regimens, even those far from the MTD-regimen. This phenomenon could be due to the additional PK/PD information and the correct understanding of the toxicity mechanism. Using the CRM, the entire dose regimen toxicity curve is well estimated when the skeleton is close to the truth, as shown in Scenario 1 (Web Appendix F). However, in most cases, the toxicity estimation is precise around the MTD-regimen, but not reliable for the other dose regimens. Regarding the dose regimens far from the MTD-regimen, the hierarchical-DRtox seems to estimate the toxicity probability with less bias but more variance than the logistic-DRtox. In Web Appendix F, the distribution of the root mean square error (RMSE) of all methods is plotted; the RMSE is computed on all dose regimens or

on the MTD-regimen and its neighbors. Near the MTD-regimen, the estimation of the logistic-DRtox is better than that of the hierarchical-DRtox; both models are better than the CRM. However, in the scenarios in which the MTD-regimen is at extreme positions of the dose regimens panel (Scenarios 2 and 4 in Web Appendix F), the entire dose regimen toxicity relationship is better estimated with the hierarchical-DRtox than the logistic-DRtox.

4.2.3 | Recommendation of a more suitable untested dose regimen

Finally, one of the strengths of the DRtox is that it models the entire relationship between the dose regimen and toxicity and can predict the toxicity probability of any new dose regimen. Notably, in this work, we assumed that the administration times were fixed to simplify the notations




FIGURE 4 Violin plots of the estimated toxicity probabilities in an additional scenario in which the dose regimen panel missed the true MTD-regimen and in Scenario 3 on 1000 trials implemented with the CRM including 30 patients. The predicted toxicity probability of a new regimen S_{new} is framed in dotted line. Horizontal lines on the density estimates represent the median and first and third quartiles of the distributions and the plus sign represents the mean. The dashed line represents the toxicity target and the solid line represents the true toxicity probabilities

but regimens with different times of drug administration can also be considered. Therefore, at the end of the dose-escalation stage of the trial, the DRtox can recommend dose regimens that were not tested in the trial to be investigated in expansion studies. For example, let us imagine a scenario in which the panel of dose regimens missed the true MTD-regimen, as illustrated in the upper plot of Figure 4, where regimen $S_3 = (5, 10, 25, 50, 50, 50, 50) \mu\text{g}/\text{kg}$ is underdosing and regimen $S_4 = (10, 25, 50, 100, 150, 150, 150) \mu\text{g}/\text{kg}$ is overdosing.

The upper plot of Figure 4 illustrates the gap between the estimated toxicity probabilities of regimens S_3 and S_4 , suggesting that an alternative regimen could be found to have a toxicity probability closer to the target. At the end of the dose-escalation stage of the trial, the DRtox can predict the toxicity probability of any new regimen, such as regimen $S_{new} = (10, 25, 50, 100, 100, 100, 100) \mu\text{g}/\text{kg}$, whereas

the CRM is unable to perform predictions as the model is built on a skeleton based on the panel of dose regimens. In the upper plot of Figure 4, we can observe that both the hierarchical-DRtox and the logistic-DRtox predict that new regimen S_{new} has a toxicity probability closer to the target; therefore we can propose to evaluate the new regimen in expansion cohorts.

Another practical case is illustrated in Scenario 3 in which the objective was to administer the steady-state dose of $40 \mu\text{g}/\text{kg}$ as soon as possible. As shown in the lower plot of Figure 4, the estimated MTD-regimen is $S_4 = (10, 20, 40, 40, 40, 40, 40) \mu\text{g}/\text{kg}$, and the next regimen of the panel, $S_5 = (20, 40, 40, 40, 40, 40, 40)$, is estimated to be too toxic. Nevertheless, one might wonder whether another regimen with an acceptable toxicity could be found in which the steady-state dose is administered from the second administration. The DRtox

predicts the toxicity probability of new regimen $S_{\text{new}} = (10, 40, 40, 40, 40, 40, 40)$ to be approximately 0.3 as shown in the lower plot of Figure 4, and this new regimen can be compared in terms of efficacy to the estimated MTD-regimen S_4 in subsequent stages of the trial. Therefore, at the end of the trial, the DRtox can evaluate alternative regimens that were not included in the panel for future studies.

5 | DISCUSSION

In this work, we developed the DRtox approach to model the relationship between the dose regimen and toxicity by modeling a PD endpoint. We estimated the toxicity related to the PD endpoint in the context of an ongoing phase I trial in which the assumption of a monotonic increase in the dose regimen toxicity relationship did not hold. We found that when the process generating toxicity was reasonably understood and approximated, adding PK/PD information increased the PCS. This method allowed for a better estimation of the dose regimen toxicity curves, as this type of modeling enabled the sharing of more information across regimens. Moreover, the DRtox was able to evaluate additional regimens for expansion cohorts that were not present in the dose regimen panel but may have a predicted toxicity probability closer to the target. In practice, our methods should be applied at the end of the dose-escalation phase of the motivating trial once all PK/PD and toxicity data are collected. Our model can address missing data as follows: (1) Regarding missing doses in the dose regimen and associated cytokine profiles, as we are using nonlinear PK/PD modeling, our method would take into account whether a patient misses one or more planned doses as the model considers the actual regimen received and not the planned regimen. (2) Regarding missing cytokine data, which is expected to be rare in this trial as the cytokine is carefully monitored by frequent sampling to detect its peak, individual cytokine peaks could be predicted from the population PK/PD model. However, it would be more common for PK/PD data to be below the limit of quantification, but these data are considered by the PK/PD model as censored data rather than missing data. (3) During the enrollment, and due to the sequential feature of the dose-allocation design, patients with missing or with nonevaluable toxicity outcome are replaced. In this case, during the enrollment, new patients are treated at the same dose-levels to account for the design requirements.

In the simulation study, we assumed that the dose regimens were ordered, but the DRtox can be applied when only partial ordering is known. As the DRtox is applied at the end of the trial, the choice of the dose-escalation design may have a significant impact on the results. The performance achieved using a model-based design, such as the

CRM with 30 patients, is better than that achieved using an algorithm-based design, such as the 3+3 design, which has the main disadvantage of treating most patients at sub-therapeutic doses and having a small total sample size that cannot be fixed before the trial.

Regarding the logistic-DRtox, since drug administration is stopped in the case of toxicity, the performance can be impacted by incomplete observations of the PD endpoint, even though it seemed to lightly impact our simulation study. In the case toxicities occur at the beginning of the administrations, resulting in a high number of incomplete PD observations, we propose the use of the predicted PD given by the PK/PD model under the complete regimen planned.

The hierarchical-DRtox added a constraint, that is, toxicity must occur at the maximum of the PD response. Errors in the PK/PD estimation may lead to an undefined hierarchical model. In our simulation study, we observed this latter issue in less than 2% of the trials. In the real world, this issue could indicate that the proposed PK/PD model is incorrect, and that another model should be considered. However, in our simulation study, we decided to run other simulated datasets for all methods to replace the 2% of the trials having the issue defined above. One way to relax this constraint is to allow the patients' toxicity threshold to vary among administrations by adding a second latent variable, which could lead to complex models that are challenging to estimate.

In this work, we assumed that all dose regimens have the same repetition scheme and duration. However, the DRtox can address regimens with different schemes, administration modes, etc. The first part of the DRtox relies on PK/PD modeling, so an incorrect PK/PD model may have a negative impact on the full method. However, as usual in the PK/PD field, the aim of the modeling is to have a good fit/prediction of the patients' PK/PD profiles even with simplified models. Therefore, an approximated PK/PD model could still be applied without DRtox performance loss if the PD endpoint is well fitted.

In conclusion, we proposed a general approach for modeling toxicity through a PK/PD endpoint. In this work, we considered a specific PD endpoint in the context of an actual ongoing clinical trial, but various endpoints (such as the AUC or a combination of several toxicity biomarkers) could be used depending on the type of toxicity considered. Moreover, we developed the DRtox under the assumption that toxicity was linked to the maximum value of the PD biomarker, but other assumptions could be raised, such as assuming a cumulative effect. The usual dose-finding designs were developed to determine the MTD in the first cycle of treatment after a single administration. However, with the increase in the number of targeted molecules, immuno-oncology therapies, and combinations with

alternative dose regimens, standard dose-allocation designs fail to identify the dose regimen recommended for future studies. Incorporating PK/PD exposure data in early phase toxicity modeling through stronger collaboration between biostatisticians and pharmacometricians may lead to a better understanding of the entire dose regimen toxicity relationship and provide alternative dosage recommendation for the next phases of the clinical development.

ACKNOWLEDGMENTS

This work was partially funded by a grant from the Association Nationale de la Recherche et de la Technologie, with Sanofi-Aventis R&D, Convention industrielle de formation par la recherche number 2018/0530. The authors would like to thank the associate editor and the two referees for their constructive comments that led to significant improvements of the manuscript. They also would like to thank Raouf El-Cheikh, Laurent Nguyen, and Christine Veyrat-Follet for their time and explanations of PK/PD modeling and Paula Fraenkel for her thoughtful review of the manuscript. Marie-Karelle Riviere and Moreno Ursino made equal contributions and are co-last authors.

DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article as no new data were created or analyzed in this paper.

ORCID

Emma Gerard  <https://orcid.org/0000-0003-2498-4801>

Sarah Zohar  <https://orcid.org/0000-0002-8429-2340>

Moreno Ursino  <https://orcid.org/0000-0002-5709-4322>

REFERENCES

- Berry, S.M., Carlin, B.P., Lee, J.J. and Muller, P. (2010) *Bayesian adaptive methods for clinical trials*. Boca Raton, FL: CRC Press.
- Boissel, N., de Botton, S., Thomas, X.G., Rao, E., Bonnevaux, H., Rubin-Carrez, C. et al. (2018) An open-label, first-in-human, dose escalation study of a novel CD3-CD123 bispecific T-cell engager administered as a single agent by intravenous infusion in patients with relapsed or refractory acute myeloid leukemia, B-cell acute lymphoblastic leukemia, or high risk myelodysplastic syndrome. *American Society of Clinical Oncology*, 36, TPS7076-TPS7076.
- Braun, T.M., Thall, P.F., Nguyen, H. and de Lima, M. (2007) Simultaneously optimizing dose and schedule of a new cytotoxic agent. *Clinical Trials*, 4, 113–124.
- Braun, T.M., Yuan, Z. and Thall, P.F. (2005) Determining a maximum-tolerated schedule of a cytotoxic agent. *Biometrics*, 61, 335–343.
- Bullock, J.M., Lin, T. and Bilic, S. (2017) Clinical pharmacology tools and evaluations to facilitate comprehensive dose finding in oncology: a continuous risk-benefit approach. *The Journal of Clinical Pharmacology*, 57, S105–S115.
- Chen, X., Kamperschroer, C., Wong, G. and Xuan, D. (2019) A modeling framework to characterize cytokine release upon T-cell-engaging bispecific antibody treatment: methodology and opportunities. *Clinical and Translational Science*, 12(6), 600–608.
- Cheung, Y.K. (2011) *Dose finding by the continual reassessment method*. Boca Raton, FL: CRC Press.
- Chevret, S. (2006) *Statistical methods for dose-finding experiments (Vol. 24)*. Chichester: Wiley.
- Crowley, J. and Hoering, A. (2012) *Handbook of statistics in clinical oncology*. Boca Raton, FL: CRC Press.
- Gelman, A. (2006) Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Analysis*, 1(3), 515–534.
- Günhan, B.K., Weber, S. and Friede, T. (2020) A Bayesian time-to-event pharmacokinetic model for phase I dose-escalation trials with multiple schedules. *Statistics in Medicine*, 39(27), 3986–4000.
- Guo, B., Li, Y. and Yuan, Y. (2016) A dose-schedule finding design for phase I–II clinical trials. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 65, 259–272.
- Li, Y., Bekele, B.N., Ji, Y. and Cook, J.D. (2008) Dose-schedule finding in phase I/II clinical trials using a Bayesian isotonic transformation. *Statistics in Medicine*, 27, 4895–4913.
- Liu, C.A. and Braun, T.M. (2009) Parametric non-mixture cure models for schedule finding of therapeutic agents. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 58, 225–236.
- Lixoft SAS (2019) Monolix version 2019R1. Antony, France. Available at: <http://lixoft.com/products/monolix/>
- Lyu, J., Curran, E. and Ji, Y. (2018) Bayesian adaptive design for finding the maximum tolerated sequence of doses in multicycle dose-finding clinical trials. *JCO Precision Oncology*, 2, 1–19.
- Morita, S., Thall, P.F. and Müller, P. (2008) Determining the effective sample size of a parametric prior. *Biometrics*, 64(2), 595–602.
- Musuamba, F.T., Manolis, E., Holford, N., Cheung, S.Y.A., Friberg, L.E., Ogungbenro, K. et al. (2017) Advanced methods for dose and regimen finding during drug development: summary of the EMA/EFPIA workshop on dose finding (London 4–5 December 2014). *CPT: Pharmacometrics and Systems Pharmacology*, 6(7), 418–429.
- O’Quigley, J., Pepe, M. and Fisher, L. (1990) Continual reassessment method: a practical design for phase 1 clinical trials in cancer. *Biometrics*, 33–48.
- R Core Team (2018) R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. Available at: <https://www.R-project.org/>.
- Schmoor, C. and Schumacher, M. (1992) Adaptive statistical procedures for the analysis of nonmonotone dose–response relationships. *Biometrie und Informatik in Medizin und Biologie*, 23, 113–126.
- Shimabukuro-Vornhagen, A., Gödel, P., Subklewe, M., Stemmler, H.J., Schlößer, H.A., Schlaak, M. et al. (2018) Cytokine release syndrome. *Journal For Immunotherapy Of Cancer*, 6, 56.
- Stan Development Team (2019) RStan: the R interface to Stan. R package version 2.19.2. Available at: <http://mc-stan.org/>.
- Storer, B.E. (1989) Design and analysis of phase I clinical trials. *Biometrics*, 45, 925–937.
- Teachey, D.T., Lacey, S.F., Shaw, P.A., Melenhorst, J.J., Maude, S.L., Frey, N. et al. (2016) Identification of predictive biomarkers for cytokine release syndrome after chimeric antigen receptor T-cell therapy for acute lymphoblastic leukemia. *Cancer Discovery*, 6(6), 664–679.
- Thall, P.F., Nguyen, H.Q., Braun, T.M. and Qazilbash, M.H. (2013) Using joint utilities of the times to response and toxicity to adaptively optimize schedule–dose regimens. *Biometrics*, 69, 673–682.
- Ursino, M., Zohar, S., Lentz, F., Alberti, C., Friede, T., Stallard, N. et al. (2017) Dose-finding methods for Phase I clinical trials using

- pharmacokinetics in small populations. *Biometrical Journal*, 59, 804–825.
- US Department of Health and Human Services (2017) Common terminology criteria for adverse events (CTCAE) v5.0.
- Wages, N.A., O'Quigley, J. and Conaway, M.R. (2014) Phase I design for completely or partially ordered treatment schedules. *Statistics in Medicine*, 33, 569–579.
- Yuan, Y., Nguyen, H.Q. and Thall, P.F. (2017) *Bayesian designs for phase I-II clinical trials*. Boca Raton, FL: CRC Press.
- Zhang, J. and Braun, T.M. (2013) A phase I Bayesian adaptive design to simultaneously optimize dose and schedule assignments both between and within patients. *Journal of the American Statistical Association*, 108, 892–901.
- Zhu, M., Wu, B., Brandl, C., Johnson, J., Wolf, A., Chow, A. et al. (2016) Blinatumomab, a bispecific T-cell engager (BiTE®) for CD-19 targeted cancer immunotherapy: clinical pharmacology and its implications. *Clinical Pharmacokinetics*, 55(10), 1271–1288.

SUPPORTING INFORMATION

Web Appendices, Tables, and Figures referenced in Sections 3 and 4 and the code are available with this paper at the Biometrics website on Wiley Online Library.

How to cite this article: Gerard E, Zohar S, Thai H-T, Lorenzato C, Riviere M-K, Ursino M. Bayesian dose regimen assessment in early phase oncology incorporating pharmacokinetics and pharmacodynamics. *Biometrics*. 2021;1–13.
<https://doi.org/10.1111/biom.13433>

Chapter 5

Bayesian modeling of a bivariate toxicity outcome for early phase oncology trials evaluating dose regimens

In this chapter, we describe the modeling of the relationship between the dose regimen and a bivariate toxicity outcome that resulted in a publication in *Statistics in Medicine* [?]. A summary is given in Section 5.1 and the complete paper is provided in Section 5.2.

5.1 Summary

Context In Chapter 4, we proposed the DRtox approach that models the probability of a PD-related toxicity (CRS) by incorporating PK/PD models in order to describe the reduction of the risk of this toxicity with intra-patient dose-escalation. This method was developed assuming that this PD-related toxicity was the only DLT that could occur. However, other DLTs, that we name DLT_o , could be observed and should be considered when recommending the MTD-regimen. Unlike the first type of toxicity, the PK/PD process generating the DLT_o might be unknown and no prior assumption could be raised to relate DLT_o with multiple drug administrations. Moreover, the PD-related toxicity and DLT_o might be associated, for example, CRS can be associated with neurotoxicities as highlighted in Section 3.1. The objective was therefore to extend the DRtox approach to include the modeling of DLT_o in order to determine the MTD-regimen accounting for all toxicity types.

Method We proposed to model the DLT as a bivariate binary outcome, to distinguish the PD-related outcome from the DLT_o , with multiple administrations of varying doses. Our modeling strategy can be divided in three steps.

In the first step, the PD-related outcome was modeled by incorporating PK/PD models using the logistic DRtox approach (as developed in Chapter 4).

In the second step, we proposed to model the cumulative probability of DLT_o , using either a marginal or a conditional formulation, and defined the cumulative dose to account for the multiple drug administrations. We chose to model the cumulative probability of DLT_o instead of the conditional probability of DLT_o at each administration because the latter was expected to be very small as we had to deal with two different types of DLT. We only considered the cumulative dose as a predictor for the model to remain as simple as possible due to the context of small sample size.

Finally, we proposed to model the DLT as a bivariate binary outcome and developed three joint approaches. The first approach, DRtox_indep, assumed independence between the two types of DLT and estimated the probability of DLT from the marginal probabilities of CRS and DLT_o . The second approach, DRtox_copula, explicitly modeled the correlation between the two toxicities using a copula model, similarly to toxicity/efficacy modeling or combination trials as developed in Sections 2.2.1.3 and 2.2.1.4. DRtox_copula modeled the probability of DLT from the marginal probabilities of each type of DLT and an additional parameter. The third approach, DRtox_cond, accounted for the association between toxicities without explicitly modeling it. Indeed, this approach modeled the joint probability of DLT from the marginal probability of the PD-related toxicity and the conditional probability of DLT_o given no CRS has occurred. This conditional modeling was motivated by the assumptions on the motivating trial that CRS tend to occur at the beginning of the administrations so CRS might prevent future DLT_o to occur (due to the remaining administrations planned) as drug administration is stopped in case a DLT occurs.

The prior distributions were elicited from the initial guesses of the probabilities of DLT, assuming that the PD-related toxicity and DLT_o were independent and had equal probabilities. The amount of information provided by the prior distributions

was measured by the prior ESS.

Results We evaluated the performance of the three joint approaches on an extensive simulation study. We applied our methods at the end of the dose-escalation stage of the trial designed with a modified CRM. We considered multiple scenarios that evaluated multiple positions of the true MTD-regimen, the proportion of each type of DLT and the association between the two types of DLT. We observed good performance in selecting the correct MTD-regimen. We also observed that DRtox_copula and DRtox_cond had a higher proportion of correct selection (PCS) than DRtox_indep when increasing the association between the two types of DLT with no loss of performance in case of true independence. Finally, our three joint approaches preserved the DRtox ability of predicting the probability of DLT of untested dose regimens.

Conclusion We proposed to model the probability of DLT with the dose regimen by distinguishing a PD-related toxicity from other DLTs, named DLT_o . This bivariate modeling allowed to account for a complex assumption on multiple administrations for the PD-related toxicity, while no such assumption could be raised for DLT_o . The three joint approaches exhibited good performance on a simulation study.

Our work has some limitations. We considered a simple model on the DLT_o due to the context of two types of toxicity and small sample size. We also observed that the parameter of the copula was difficult to be estimated even in case of high association so the simple independent model or the conditional model might be preferred. We also observed that DRtox_indep and DRtox_copula had a tendency to underestimate the marginal probability of DLT_o due to the fact that DLT_o might not be observed in case of CRS, which motivated the DRtox_cond approach. Finally, our modeling strategy requires multiple steps of modeling so it should be applied when enough data is gathered but it can help support the choice of the dose regimen for the future development of the molecule using additional information collected in early phase trials.

5.2 Publication

The supporting information of this paper is provided in Appendix D.

RESEARCH ARTICLE

Bayesian modeling of a bivariate toxicity outcome for early phase oncology trials evaluating dose regimens

Emma Gerard^{1,2,3,4} | Sarah Zohar^{1,2} | Christelle Lorenzato³ |
Moreno Ursino^{1,2,5} | Marie-Karelle Riviere⁴

¹Inserm, Centre de Recherche des Cordeliers, Université de Paris, Sorbonne Université, Paris, France

²HeKA, Inria, Paris, France

³Oncology Biostatistics, Biostatistics and Programming Department, Sanofi R&D, Vitry-sur-Seine, France

⁴Statistical Methodology Group, Biostatistics and Programming Department, Sanofi R&D, Chilly-Mazarin, France

⁵Unit of Clinical Epidemiology, AP-HP, CHU Robert Debré, Université de Paris, Sorbonne Paris-Cité, Inserm U1123 and CIC-EC 1426, Paris, France

Correspondence

Sarah Zohar, Inserm, Centre de Recherche des Cordeliers, Université de Paris, Sorbonne Université, 15 rue de l'École de Médecine, 75006 Paris, France.
Email: sarah.zohar@inserm.fr

Most phase I trials in oncology aim to find the maximum tolerated dose (MTD) based on the occurrence of dose limiting toxicities (DLT). Evaluating the schedule of administration in addition to the dose may improve drug tolerance. Moreover, for some molecules, a bivariate toxicity endpoint may be more appropriate than a single endpoint. However, standard dose-finding designs do not account for multiple dose regimens and bivariate toxicity endpoint within the same design. In this context, following a phase I motivating trial, we proposed modeling the first type of DLT, cytokine release syndrome, with the entire dose regimen using pharmacokinetics and pharmacodynamics (PK/PD), whereas the other DLT (DLT₀) was modeled with the cumulative dose. We developed three approaches to model the joint distribution of DLT, defining it as a bivariate binary outcome from the two toxicity types, under various assumptions about the correlation between toxicities: an independent model, a copula model and a conditional model. Our Bayesian approaches were developed to be applied at the end of the dose-allocation stage of the trial, once all data, including PK/PD measurements, were available. The approaches were evaluated through an extensive simulation study that showed that they can improve the performance of selecting the true MTD-regimen compared to the recommendation of the dose-allocation method implemented. Our joint approaches can also predict the DLT probabilities of new dose regimens that were not tested in the study and could be investigated in further stages of the trial.

KEYWORDS

Bayesian joint modeling, bivariate toxicity, cumulative probability of toxicity, dose regimen, early phase oncology, pharmacokinetics/pharmacodynamics

1 | INTRODUCTION

Most phase I dose-finding trials in oncology aim to determine the maximum tolerated dose (MTD), which is defined as the highest dose that does not exceed a predefined probability of dose-limiting toxicity (DLT), in a prespecified observational window. The DLT is a binary outcome defined to summarize the patient's toxicity profile and is usually derived from

Moreno Ursino and Marie-Karelle Riviere contributed equally to this study.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2021 The Authors. *Statistics in Medicine* published by John Wiley & Sons Ltd.

toxicity classification in severity grades according to the National Cancer Institute (NCI) Common Toxicity Criteria for Adverse Events.¹ For instance, a DLT may be defined as the occurrence of a grade 3 or higher toxicity observed in at least one organ. Algorithm-based designs, such as the 3+3,² or model-based designs, such as the continual reassessment method (CRM),³ are the most common adaptive approaches for recommending the MTD based on the occurrence of DLT. Evaluating the schedule of administration, in addition to the dose, is increasingly being investigated during this phase to improve drug tolerance. For example, lead-in dose step-up dosing is common in immunotherapy trials.⁴⁻⁸ However, traditional designs reduce the entire dose regimen, defined by the NCI as “a treatment plan that specifies the dosage, the schedule, and the duration of treatment”⁹, to a single value.

While a binary toxicity outcome is convenient in the context of small sample size, various types of DLT might not be identified as equally important.¹⁰ Several authors have proposed to define a continuous toxicity score to differentiate the types of toxicity¹⁰⁻¹² and to define multiple toxicity constraints on the toxicity types and grades to define the MTD.¹³⁻¹⁶

Several methods have been developed to consider more complex schedules of administration. Dose-finding designs have been proposed to evaluate the dose-schedule combination in a two-dimensional framework.^{17,18} The toxicity outcome can also be modeled in a longitudinal setting where a dose regimen composed of multiple administrations of the same or various doses is evaluated. Legedza and Ibrahim¹⁹ proposed modeling the repeated treatment administrations using the total amount of drug in the body based on pharmacokinetic principles. Other methods account for the cumulative effect of the dose regimens using a time to toxicity outcome and by modeling the total hazard of toxicity for the dose regimen as the sum of the hazard of each administration.²⁰⁻²³ Another possibility, proposed by Fernandes et al,²⁴ is to estimate the conditional probability of toxicity at each administration given that no toxicity was observed in the previous administrations. Ursino et al²⁵ developed a model for the cumulative probability of toxicity by accounting for the first dose administered and the cumulative dose. Finally, to model the cumulative effect of multiple administrations, other authors have included pharmacokinetic and pharmacodynamic (PK/PD) modeling in the dose regimen recommendation. Günhan et al²⁶ considered latent PK profiles to measure drug exposure in a Bayesian adaptive model, and Gerard et al proposed complete PK/PD modeling to assess the relationship between the dose regimen and a specific toxicity outcome.

Motivating trial

This work is based on an ongoing first-in-human dose-escalation trial of SAR440234²⁷ administered as a single agent to patients with relapsed or refractory acute myeloid leukemia, high-risk myelodysplastic syndrome or B-cell acute lymphoblastic leukemia (NCT03594955²⁸). SAR440234 is a novel bispecific T-cell engager antibody that activates and redirects cytotoxic T lymphocytes (CTLs) to enhance the CTL-mediated elimination of CD123-expressing tumor cells. Due to this mechanism of action, the main toxicity expected in this trial is the cytokine release syndrome (CRS), which is a systemic inflammatory response that has been associated with the peak of cytokine, considered as a pharmacodynamic endpoint.^{29,30} Since repeated administrations of the drug can reduce the risk of CRS,³¹ dose regimens with a fixed inpatient dose-escalation scheme are evaluated in this trial as already proposed for other similar molecules.⁴⁻⁸ Namely, before administering repeatedly the steady-state dose, defined as the maximum dose the regimen should reach, patients are treated with lower initial increasing doses, defined as the lead-in doses. This administration scheme attempts to decrease the peak of cytokine observed when the steady-state dose is administered directly. The objective of the trial is to find the MTD of SAR440234, and a standard dose-escalation procedure is implemented (3+3 design), in which each dose regimen is associated with a single dose-level.³²

Objective

To obtain a better estimate of the maximum tolerated dose regimen (MTD-regimen), defined as the dose regimen having the probability of DLT closest to the toxicity target, at the end of the trial, a Bayesian dose regimen assessment method (DRtox) was firstly developed.³³ The DRtox proposed to model the binary CRS incorporating pharmacokinetic and pharmacodynamic (PK/PD) modeling. Since the development of the DRtox, the investigators have gained knowledge about the potential occurrences of other important dose-limiting toxicities (DLT_o) that should be considered when recommending the MTD-regimen. Unlike the CRS, the PK/PD process generating the DLT_o might be unknown.³⁴ Moreover, the CRS can be associated with several types of other toxicities.³⁵⁻³⁸ Thus, a new modeling approach that accounts for the potential interactions between the CRS and the DLT_o and the multiple dose administrations was required. Since the CRS and

DLT_0 are clinically different and can be distinguished, we propose to model the DLT as a bivariate binary endpoint with the dose regimen, preserving the PK/PD approach for the CRS and adding a cumulative modeling approach for the DLT_0 .

The objective of this paper is to model the probability of DLT as a bivariate binary endpoint (CRS and DLT_0) in the context of multiple dose administrations in order to determine the MTD-regimen while preserving the modeling of the CRS with PK/PD. As PK/PD measures are usually analyzed by batch and may therefore not be available in real time after each cohort is treated, a complete sequential dose-allocation approach incorporating PK/PD can be difficult to perform in practice. We then develop our modeling strategy to be performed at the end of the trial, that is, once all data have been collected.

2 | MODEL

Let $\mathcal{D} = \{d_1, \dots, d_L\}$ be the set of doses that can be administered to patients and let $\mathcal{S} = \{\mathbf{S}_1, \dots, \mathbf{S}_K\} \subset \mathbb{S}$ be the panel of K dose regimens to be studied in the trial. The dose regimen $\mathbf{S}_k \in \mathcal{S}$ is defined as a sequence of J doses, $\mathbf{S}_k = (d_{k,1}, d_{k,2}, \dots, d_{k,J})$, administered at times $\mathbf{t} = (t_1, t_2, \dots, t_J)$, where $d_{k,j} \in \mathcal{D}$, $j \in \{1, \dots, J\}$ and $k \in \{1, \dots, K\}$. Let $n \in \mathbb{N}$ be the number of patients included in the trial. Let $Y_{i,j}^{(1)}$ be the binary CRS response and $Y_{i,j}^{(2)}$ be the binary DLT_0 response (another DLT, different from the CRS) of patient i observed exactly after the j^{th} administration. As we assume that both a CRS and a DLT_0 can occur during the same administration, the binary DLT response $Y_{i,j}$ is defined as $Y_{i,j} = \max(Y_{i,j}^{(1)}, Y_{i,j}^{(2)})$. Let Y_i , $Y_i^{(1)}$, and $Y_i^{(2)}$ be the global DLT, CRS, and DLT_0 responses of patient i at the end of the dose regimen, respectively. Let $\tilde{\mathbf{s}}_i = (d_{i,1}, d_{i,2}, \dots, d_{i,J}) \in \mathcal{S}$ be the dose regimen planned for the i^{th} patient. As we assume that the administration of the drug is stopped if a DLT occurs (either a CRS or a DLT_0), let j_i be the final administration of patient i . Let $\mathbf{s}_i = (d_{i,1}, d_{i,2}, \dots, d_{i,j_i}) \subset \tilde{\mathbf{s}}_i$ be the actual dose regimen received by patient i ; this implies that $\mathbf{s}_i = \tilde{\mathbf{s}}_i$ if no DLT occurs. The aim is to estimate the MTD-regimen at the end of the trial, and this is defined as the dose regimen with the probability of DLT closest to the target toxicity rate δ_T . To define the prior distributions, let (π_1, \dots, π_K) be the initial guesses of the DLT probabilities of dose regimens $(\mathbf{S}_1, \dots, \mathbf{S}_K)$, where a selected probability, π_{k_T} , is equal to the target δ_T .

The proposed modeling strategy is divided into three steps, as shown in Figure 1. In the first step, the CRS is modeled via PK/PD outcomes using the DRtox approach.³³ In the second step, the DLT_0 , using either a marginal or conditional definition, is modeled via the cumulative dose to account for the multiple dose administrations. Finally, in the third step, the DLT is modeled as a bivariate binary endpoint using both previous models under three approaches: the DRtox_indep assumes independence between the two types of toxicity, the DRtox_copula models the correlation between them via a

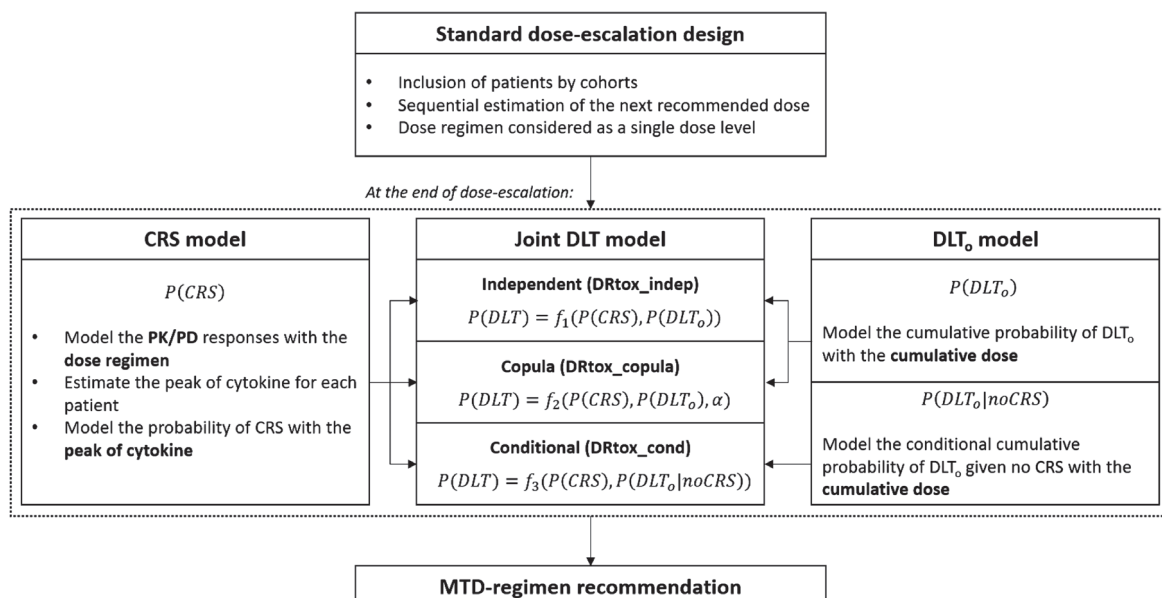


FIGURE 1 Illustration of the proposed modeling process of the bivariate toxicity endpoint at the end of the dose-escalation phase of the trial

copula model and the DRtox_cond accounts for the potential correlation between them via a conditional modeling. In the following sections, the three steps are detailed.

2.1 | CRS modeling

As the CRS is assumed to be linked to the peak of cytokine, considered a PD endpoint, we propose to model the binary CRS using the logistic-DRtox approach³³ that incorporates PK/PD modeling. The logistic-DRtox approach is performed in three steps:

1. The continuous concentration response (PK) and cytokine response (PD) are linked to the dose regimen using non-linear mixed effects models. The peak of cytokine r_i^M is predicted for each patient i according to the received dose regimen \mathbf{s}_i .
2. The predicted peak of cytokine is used to model the binary CRS via the following Bayesian model:

$$\text{logit} \left(\mathbb{P} \left(Y_i^{(1)} = 1 \right) \right) = \beta_{0,1} + \beta_{1,1} \log \left(\frac{r_i^M}{\bar{r}_{k_T}^M} \right) \quad (1)$$

where $\beta_{1,1} > 0$ is required for the probability of CRS to increase with the peak of cytokine. We normalize the peak of cytokine for prior elicitation using $\bar{r}_{k_T}^M$, which is the population value of the peak of cytokine of dose regimen \mathbf{S}_{k_T} that we initially guess to have a DLT probability of δ_T .

3. The posterior marginal CRS probability $p_k^{(1)}$ of dose regimen \mathbf{S}_k is estimated by a Monte-Carlo method to integrate both the PK/PD estimated variability and the posterior distributions resulting from the Bayesian logistic model. A point estimate is obtained by the posterior mean.

Regarding prior distributions, we consider a normal distribution for the intercept, $\beta_{0,1} \sim \mathcal{N} \left(\bar{\beta}_{0,1}, \sigma_{\beta_{0,1}}^2 \right)$, and a gamma distribution for the slope to ensure positivity, $\beta_{1,1} \sim \gamma \left(\alpha_1, \frac{\alpha_1}{\bar{\beta}_{1,1}} \right)$, where α_1 is the shape parameter, $\bar{\beta}_{0,1} = \mathbb{E}[\beta_{0,1}]$ and $\bar{\beta}_{1,1} = \mathbb{E}[\beta_{1,1}]$. The values of $\bar{\beta}_{0,1}$ and $\bar{\beta}_{1,1}$ are elicited from the prior guesses of the probability of CRS. Additional details on prior elicitation can be found in Appendix A.

2.2 | DLT₀ modeling

To perform the joint DLT modeling, we model either the marginal probability of DLT₀ (a DLT other than a CRS) or the conditional probability of DLT₀ given that no CRS has occurred. Both approaches are based on the cumulative probability of DLT₀.

2.2.1 | Marginal DLT₀ modeling

Contrary to the CRS, we do not assume that the DLT₀ is linked to a PD endpoint. To account for the multiple dose administrations that constitute the dose regimen, we model the cumulative probability of DLT₀²⁵ for patient i at administration j , and this is defined as:

$$p_{i,j}^{(2)\text{cum}} = \mathbb{P} \left(\exists! l \in \{1, \dots, j\} Y_{i,l}^{(2)} = 1 \right) = \sum_{l=1}^j \mathbb{P} \left(Y_{i,1}^{(2)} = 0, Y_{i,2}^{(2)} = 0, \dots, Y_{i,l-1}^{(2)} = 0, Y_{i,l}^{(2)} = 1 \right). \quad (2)$$

The likelihood is computed as $\prod_{i=1}^n \left(p_{i,j_i}^{(2)\text{cum}} - p_{i,j_i-1}^{(2)\text{cum}} \right)^{Y_i^{(2)}} \left(1 - p_{i,j_i}^{(2)\text{cum}} \right)^{1-Y_i^{(2)}}$ from the following probabilities of experiencing a DLT₀ and not experiencing a DLT₀ at administration j :

$$\begin{cases} \mathbb{P} \left(Y_{i,1}^{(2)} = 0, Y_{i,2}^{(2)} = 0, \dots, Y_{i,j-1}^{(2)} = 0, Y_{i,j}^{(2)} = 1 \right) = p_{i,j}^{(2)\text{cum}} - p_{i,j-1}^{(2)\text{cum}} \\ \mathbb{P} \left(Y_{i,1}^{(2)} = 0, Y_{i,2}^{(2)} = 0, \dots, Y_{i,j-1}^{(2)} = 0, Y_{i,j}^{(2)} = 0 \right) = 1 - p_{i,j}^{(2)\text{cum}} \end{cases}$$

Since we model two types of toxicity but with a standard joint target toxicity, we have to deal with very small marginal probabilities of toxicity. Therefore, we propose to adapt the model proposed by Ursino et al²⁵ for inpatient dose-escalation. As DLT₀ is only one type of the bivariate DLT, the number of observed DLT₀ is relatively small; thus we propose to reduce the number of parameters to two as follows:

$$\text{logit} \left(p_{i,j}^{(2)\text{cum}} \right) = \beta_{0,2} + \exp(\beta_{1,2}) \log \left(\frac{1}{D_{k_T}} \sum_{l=1}^j d_{i,l} \right) \quad (3)$$

where $D_{k_T} = \sum_{l=1}^j d_{k_T,l}$ is the cumulative dose of the reference dose regimen, which we assume has a DLT probability of δ_T . The posterior marginal probability of DLT₀ of dose regimen \mathbf{S}_k is then defined as $p_k^{(2)} = \text{logit}^{-1} \left(\beta_{0,2} + \exp(\beta_{1,2}) \log \left(\frac{1}{D_{k_T}} \sum_{l=1}^j d_{i,l} \right) \right)$. A point estimate is obtained by the mean of the posterior toxicity distribution.

Regarding prior distributions, we consider a normal distribution for the intercept $\beta_{0,2} \sim \mathcal{N}(\bar{\beta}_{0,2}, \sigma_{\beta_{0,2}}^2)$ and the slope $\beta_{1,2} \sim \mathcal{N}(\bar{\beta}_{1,2}, \sigma_{\beta_{1,2}}^2)$, $\bar{\beta}_{0,2} = \mathbb{E}[\beta_{0,2}]$ and $\bar{\beta}_{1,2} = \mathbb{E}[\beta_{1,2}]$. Values of $\bar{\beta}_{0,2}$ and $\bar{\beta}_{1,2}$ are elicited from the prior guesses of the probability of DLT₀. Additional details on prior elicitation can be found in Appendix A.

2.2.2 | Conditional DLT₀ modeling

The conditional DLT₀ model is very similar to the marginal DLT₀ model. We define the conditional cumulative probability of DLT₀, given that no CRS has occurred, as follows:

$$p_{i,j\star}^{(2)\text{cum}} = \mathbb{P} \left(\exists! l \in \{1, \dots, j\} Y_{i,l}^{(2)} = 1 \mid Y_i^{(1)} = 0 \right). \quad (4)$$

The conditional cumulative model is applied on a subset of the population, that is, on the patients who do not experience a CRS. The likelihood is computed similarly to that of the marginal DLT₀ model as

$$\prod_{i=1}^n \mathbb{1}_{Y_i^{(1)}=0} \left(p_{i,j\star}^{(2)\text{cum}} - p_{i,j-1\star}^{(2)\text{cum}} \right)^{y_i^{(2)}} \left(1 - p_{i,j\star}^{(2)\text{cum}} \right)^{1-y_i^{(2)}}.$$

We consider the same model as the one proposed in Section 2.2.1, and it is defined as:

$$\text{logit} \left(p_{i,j\star}^{(2)\text{cum}} \right) = \beta_{0,2\star} + \exp(\beta_{1,2\star}) \log \left(\frac{\sum_{l=1}^j d_{i,l}}{D_{k_T}} \right). \quad (5)$$

The posterior conditional probability of DLT₀ given that no CRS has occurred of dose regimen \mathbf{S}_k is then defined as $p_{k\star}^{(2)} = \text{logit}^{-1} \left(\beta_{0,2\star} + \exp(\beta_{1,2\star}) \log \left(\frac{\sum_{l=1}^j d_{k,l}}{D_{k_T}} \right) \right)$. A point estimate is obtained by the mean of the posterior toxicity distribution.

We consider the same prior distributions as those of the marginal model described in Section 2.2.1.

2.3 | Joint DLT modeling

The final step is to model the probability of DLT from the joint probability of CRS and DLT₀. The probability of DLT of dose regimen \mathbf{S}_k is defined as:

$$p_k = \mathbb{P} \left(Y = 1 \mid \mathbf{S}_k \right) = 1 - \mathbb{P} \left(Y = 0 \mid \mathbf{S}_k \right) = 1 - \mathbb{P} \left(Y^{(1)} = 0, Y^{(2)} = 0 \mid \mathbf{S}_k \right). \quad (6)$$

Different assumptions can be used when modeling the joint probability of CRS and DLT₀. The easiest assumption is to consider that both toxicities are independent (conditionally on the dose regimen). However, we can also assume a correlation between the two toxicities;³⁵⁻³⁸ for example, a patient experiencing one kind of toxicity may be more sensitive to the drug than others and therefore be at higher risk of experiencing other toxicities. We propose three approaches

to compute the probability of DLT starting from the probabilities of CRS and DLT_0 using a marginal or conditional formulation.

2.3.1 | Independent model: DRtox_indep

The first assumption that can be raised is that $Y^{(1)}$ and $Y^{(2)}$ are independent (conditionally on the dose regimen \mathbf{S}_k). In this case, the posterior distribution of DLT can be directly expressed from the posterior marginal probabilities of CRS and DLT_0 defined in Sections 2.1 and 2.2.1 as follows:

$$p_k = 1 - \mathbb{P}\left(Y^{(1)} = 0 \mid \mathbf{S}_k\right) \mathbb{P}\left(Y^{(2)} = 0 \mid \mathbf{S}_k\right) = 1 - \left(1 - p_k^{(1)}\right) \left(1 - p_k^{(2)}\right). \quad (7)$$

The probability of DLT p_k is then estimated as the mean of the posterior distribution.

2.3.2 | Copula model: DRtox_copula

When dependence between both toxicities is assumed, the joint probability of DLT can be modeled via a copula distribution defined from the marginal distribution of each toxicity. In this case, the modeling process is performed in two steps. First, marginal models are developed for the CRS and the DLT_0 , as shown in Sections 2.1 and 2.2.1. Then, the copula model is fitted on all patients from the resulting point estimates of the marginal probabilities of each toxicity of all dose regimens estimated from the two previous models.

We consider the popular class of Archimedean copulas³⁹ that have already been used in phase I trials when modeling both toxicity and efficacy⁴⁰⁻⁴² or in combination trials of two agents.^{43,44} The Archimedean copula has an explicit distribution defined from the marginal distributions of each variable and an association parameter as follows:

$$C_\alpha\left(p_k^{(1)}, p_k^{(2)}\right) = \Phi_\alpha^{-1}\left(\Phi_\alpha\left(p_k^{(1)}\right) + \Phi_\alpha\left(p_k^{(2)}\right)\right), \quad (8)$$

where Φ_α is a continuous, strictly decreasing and convex function and Φ_α^{-1} is its inverse.

Let k_i be the level of the dose regimen administered to patient i . As we assume that both toxicities can be distinguished, we can define the binary bivariate model in terms of the four possible outcomes as follows⁴⁵:

$$\begin{cases} p_i^{11} = \mathbb{P}\left(Y_i^{(1)} = 1, Y_i^{(2)} = 1\right) = C_\alpha\left(p_{k_i}^{(1)}, p_{k_i}^{(2)}\right) \\ p_i^{01} = \mathbb{P}\left(Y_i^{(1)} = 0, Y_i^{(2)} = 1\right) = p_{k_i}^{(2)} - C_\alpha\left(p_{k_i}^{(1)}, p_{k_i}^{(2)}\right) \\ p_i^{10} = \mathbb{P}\left(Y_i^{(1)} = 1, Y_i^{(2)} = 0\right) = p_{k_i}^{(1)} - C_\alpha\left(p_{k_i}^{(1)}, p_{k_i}^{(2)}\right) \\ p_i^{00} = \mathbb{P}\left(Y_i^{(1)} = 0, Y_i^{(2)} = 0\right) = 1 - p_{k_i}^{(1)} - p_{k_i}^{(2)} + C_\alpha\left(p_{k_i}^{(1)}, p_{k_i}^{(2)}\right) \end{cases} \quad (9)$$

The log-likelihood is then computed as:

$$\sum_{i=1}^n y_i^{(1)} y_i^{(2)} \log\left(p_i^{11}\right) + y_i^{(1)} \left(1 - y_i^{(2)}\right) \log\left(p_i^{10}\right) + \left(1 - y_i^{(1)}\right) y_i^{(2)} \log\left(p_i^{01}\right) + \left(1 - y_i^{(1)}\right) \left(1 - y_i^{(2)}\right) \log\left(p_i^{00}\right). \quad (10)$$

The toxicity probability of dose regimen \mathbf{S}_k is then defined as $p_k = p_k^{(1)} + p_k^{(2)} - C_\alpha\left(p_k^{(1)}, p_k^{(2)}\right)$. We consider the Clayton distribution, which is defined as $C_\alpha\left(p_k^{(1)}, p_k^{(2)}\right) = \left(\max\left(p_k^{(1)-\gamma} + p_k^{(2)-\gamma} - 1, 0\right)\right)^{-1/\gamma}$, where γ is the association parameter. We assume a positive association between both toxicities and therefore consider a gamma distribution for the prior on γ . Another popular choice for the copula distribution is the Farlie-Gumbel-Morgenstern distribution (Web Appendix B.3 of Data S1). The probability of DLT p_k is estimated as the mean of the posterior distribution.

2.3.3 | Conditional model: DRtox_cond

In this approach, we assume an association between both toxicities and express the joint distribution of DLT using a conditional formulation. As we assume that the CRS has a tendency to occur at the beginning of the dose regimen while the DLT_0 occurs at the end, we express the posterior distribution of DLT from the posterior marginal probability of CRS (Section 2.1) and the posterior conditional probability of DLT_0 (Section 2.2.2) given that no CRS has occurred, as follows:

$$p_k = 1 - \mathbb{P}\left(Y^{(1)} = 0 \mid \mathbf{S}_k\right) \mathbb{P}\left(Y^{(2)} = 0 \mid Y^{(1)} = 0, \mathbf{S}_k\right) = 1 - \left(1 - p_k^{(1)}\right) \left(1 - p_{k^*}^{(2)}\right). \quad (11)$$

The probability of DLT p_k is then estimated as the mean of the posterior distribution.

3 | SIMULATION STUDY

3.1 | Simulation settings

We defined two sets of dose regimens based on the motivating trial, Set A and Set B, as shown in Table 1. Each set included six dose regimens, where each dose regimen was defined with intra-patient dose-escalation. Each dose regimen was defined as a sequence of seven doses administered at days ($t_1 = 1, t_2 = 5, t_3 = 9, t_4 = 13, t_5 = 17, t_6 = 21, t_7 = 25$) where the dose was increased during the first four administrations to reach the steady-state dose, that was repeated for the last three administrations. In both sets, we assumed that the six dose regimens were completely ordered in terms of toxicity as the doses administered in the dose regimens were increased when increasing the level of the dose regimen. Set A and Set B shared common dose regimens ($\mathbf{S}_1, \mathbf{S}_5$ and \mathbf{S}_6), but Set B included an accelerated component since the middle dose regimens (from \mathbf{S}_2 to \mathbf{S}_4) reached higher steady-state doses more quickly than in Set A. Simulations under another set of dose regimens (Set C) that differs from the motivating trial are available in the supporting information (Web Appendix C of Data S1). The dose regimen escalation in Set C was slower than in Set A and Set B because the higher the steady-state dose is, the slower it is reached.

Both toxicities, CRS and DLT_0 , were simulated under different assumptions. The CRS was simulated from the cytokine PD profile as proposed by Gerard et al.³³ The PK/PD models used for CRS simulation were inspired by the published models on blinatumomab, which is another bispecific T-cell engager that binds to CD3 on T-cells and to CD19 on tumor cells.^{31,46} To simulate the CRS from the cytokine profile, we defined a threshold τ_T for the cytokine response and assumed

TABLE 1 Dose regimens defined in Set A and Set B used in the simulation study (in $\mu\text{g}/\text{kg}$) where each dose regimen \mathbf{S}_k is defined as the sequence of seven doses administered at days ($t_1=1, t_2=5, t_3=9, t_4=13, t_5=17, t_6=21, t_7=25$)

		t_1	t_2	t_3	t_4	t_5	t_6	t_7
Set A	\mathbf{S}_1	1	5	10	20	20	20	20
	\mathbf{S}_2	1	5	10	25	25	25	25
	\mathbf{S}_3	1	5	10	30	30	30	30
	\mathbf{S}_4	1	5	10	45	45	45	45
	\mathbf{S}_5	5	10	25	75	75	75	75
	\mathbf{S}_6	10	25	50	100	100	100	100
Set B	\mathbf{S}_1	1	5	10	20	20	20	20
	\mathbf{S}_2	1	5	10	30	30	30	30
	\mathbf{S}_3	1	5	10	40	40	40	40
	\mathbf{S}_4	1	5	10	50	50	50	50
	\mathbf{S}_5	5	10	25	75	75	75	75
	\mathbf{S}_6	10	25	50	100	100	100	100

that a CRS occurred if this threshold was exceeded, that is, $Y_{ij}^{(1)} = 1$ if $\alpha_i r_{ij} \geq \tau_T$, where α_i is a log-normally distributed measure of subject sensitivity. The true probability of CRS, $p_T^{(1)}(\mathbf{S}_k)$, of each dose regimen \mathbf{S}_k was computed using a Monte-Carlo method. Additional details on the PK/PD simulation and estimation can be found in Appendix B.1.

To create an association between the CRS and the DLT₀, we simulated the DLT₀ conditionally on the CRS status. For example, if a patient experiences a CRS, s/he may be more sensitive to the drug and, consequently, would have a higher probability of DLT₀ than a patient without CRS. We therefore simulated the DLT₀ from the conditional probability of DLT₀ at each administration (1) given that no DLT₀ had occurred in the previous administrations and (2) given that the patient would experience a CRS at some point. This conditional probability was inspired by the model proposed by Fernandes et al to account for the current and previous drug administrations as follows:

$$\text{logit} \left(\mathbb{P} \left(Y_{ij}^{(2)} = 1 \mid \left(Y_{i,j-1}^{(2)} = 0, \dots, Y_{i,1}^{(2)} = 0 \right); Y_i^{(1)} = 1 \right) \right) = a + b \log(d_{ij}) + c \log(D_{i,j-1} + 1), \quad (12)$$

where d_{ij} represents the dose given to patient i during the j^{th} administration and $D_{ij} = \sum_{l=1}^j d_{il}$ is the cumulative dose using the convention $D_{i,0} = 0$. Let $\theta = (a, b, c)$ be the simulation parameters and $f_\theta(d_j, D_{j-1}) = \text{logit}^{-1}(a + b \log(d_j) + c \log(D_{j-1} + 1))$ be the simulation model.

The conditional probability of DLT₀ for patient i given that s/he would not experience CRS was defined from the previous model as follows:

$$\begin{aligned} \mathbb{P} \left(Y_{ij}^{(2)} = 1 \mid \left(Y_{i,j-1}^{(2)} = 0, \dots, Y_{i,1}^{(2)} = 0 \right); Y_i^{(1)} = 0 \right) &= \frac{1}{\lambda} \mathbb{P} \left(Y_{ij}^{(2)} = 1 \mid \left(Y_{i,j-1}^{(2)} = 0, \dots, Y_{i,1}^{(2)} = 0 \right); Y_i^{(1)} = 1 \right) \\ &= \frac{1}{\lambda} f_\theta(d_{i,j}, D_{i,j-1}), \end{aligned} \quad (13)$$

where λ represents the risk ratio of experiencing a DLT₀ during administration j (given that no DLT₀ was experienced before) for patients who would experience a CRS vs patients with no CRS. A value of $\lambda = 1$ denotes independence between the CRS and the DLT₀, $\lambda > 1$ represents positive association and $\lambda < 1$ represents negative association.

The marginal probability of DLT₀ was computed from the previous conditional probabilities and from the marginal probability of CRS as follows:

$$p_T^{(2)}(\mathbf{S}_k) = 1 - \prod_{j=1}^J (1 - f_\theta(d_{k,j}, D_{k,j-1})) p_T^{(1)}(\mathbf{S}_k) - \prod_{j=1}^J \left(1 - \frac{1}{\lambda} f_\theta(d_{k,j}, D_{k,j-1}) \right) (1 - p_T^{(1)}(\mathbf{S}_k)). \quad (14)$$

Finally, the joint probability of DLT of dose regimen \mathbf{S}_k was computed as follows:

$$p_T(\mathbf{S}_k) = 1 - \prod_{j=1}^J \left(1 - \frac{1}{\lambda} f_\theta(d_{k,j}, D_{k,j-1}) \right) (1 - p_T^{(1)}(\mathbf{S}_k)). \quad (15)$$

Details of the mathematical development are provided in Appendix B.2.

We studied six main dose regimen toxicity scenarios, which are represented in Figure 2. Each toxicity scenario was defined by the marginal probabilities of DLT, CRS, and DLT₀ and the conditional probabilities of DLT₀ given CRS and DLT₀ given no CRS. The target probability of DLT was defined as $\delta_T = 0.30$. Scenarios 1 to 5 were built on Set A of dose regimens while Scenario 6 was built on Set B. The six toxicity scenarios studied various distributions of the probabilities of DLT, CRS, and DLT₀ and various associations between the CRS and the DLT₀. The latter association was measured for each scenario as the average risk ratio over the set of dose regimens of experiencing a DLT₀ for patients who would experience a CRS versus those who would not experience a CRS. In Scenario 1, the true MTD-regimen was \mathbf{S}_4 , which had similar probabilities of CRS and DLT₀. The CRS and DLT₀ were positively correlated with an average risk ratio of 1.85, meaning that, on average, a patient experiencing a CRS has 1.85 times greater risk of also experiencing a DLT₀ than that of a patient not experiencing a CRS. In Scenario 2, the association between the CRS and the DLT₀ was increased to an average risk ratio of 5.91. In Scenarios 3 and 4, the true MTD-regimen remained \mathbf{S}_4 , but the proportion of each type of toxicity varied with a higher probability of DLT₀ and CRS in Scenarios 3 and 5, respectively. Finally, the MTD-regimen changed to dose regimens \mathbf{S}_6 and \mathbf{S}_2 for Scenarios 5 and 6, respectively. Additional scenarios are available in the supporting information (Web Appendix A.3 of Data S1).




FIGURE 2 Definitions of the six toxicity scenarios in terms of the probabilities of dose limiting toxicity (DLT), cytokine release syndrome (CRS), and other DLT (DLT₀). For each scenario, the marginal probabilities of DLT, CRS, and DLT₀ of each dose regimen are represented as solid lines, while the conditional probabilities of DLT₀ given CRS and no CRS are represented as dotted lines. The target probability of DLT is represented by a dashed horizontal line

For each scenario, we first simulated 1000 trials with a modified CRM³ using a two-parameter logistic regression model based on a skeleton for the dose-allocation rule. The skeleton, that reduced the dose regimen into a single dose level, was defined from the initial guesses of the probabilities of DLT that were set to (0.06, 0.12, 0.20, 0.30, 0.40, 0.50), leading to $k_T = 4$ for the dose regimen of reference. The CRM sequentially included 30 patients with cohorts of size 3 and recommended a MTD-regimen (simplified to a dose-level) at the end of the dose-escalation stage. We then computed our joint modeling approaches at the end of each simulated trial using all available information to recommend the MTD-regimen.

Regarding the prior distributions, we evaluated the amount of information provided by the prior through the effective sample size (ESS) which represents the number of “hypothetical patients” used to define the prior distribution.⁴⁷ For the CRS model, we considered $\sigma_{\beta_{0,1}} = 2$ and $\alpha = 5$, for the DLT₀ model, we considered $\sigma_{\beta_{0,2}} = 2$ and $\sigma_{\beta_{1,2}} = 1$, and for the Clayton copula, we considered $\gamma \sim \gamma(1, 1)$ leading to an approximated ESS of 2. Details on prior elicitation are provided in Appendix A and results with various ESS are provided in Web Appendix B.1 of Data S1.

All simulations were performed in the R environment⁴⁸ using Monolix software⁴⁹ for PK/PD estimation and Stan⁵⁰ for Bayesian analysis.

3.2 | Simulation results

We evaluated the performance of each of our three proposed joint approaches using the proportion of correct selection (PCS), the estimation of the dose-toxicity curve and the ability of the methods to predict the toxicity probability of new dose regimens.

3.2.1 | PCS of the MTD-regimen

We compared the results of the DRtox_indep, DRtox_copula, and DRtox_cond with those of the logistic CRM. The performance of each method was evaluated with respect to the PCS, which is defined as the proportion over the 1000 trials for which the correct MTD-regimen was selected. The proportions of selecting each dose regimen in all scenarios are shown in Table 2. The results of additional scenarios can be found in Web Appendix A.3 of Data S1.

In Scenarios 1 to 4, where the true MTD-regimen is S_4 , the DRtox_indep, DRtox_copula, and DRtox_cond outperform the CRM by 10% in terms of the PCS. All three joint methods achieve very similar results and are robust to an unbalanced distribution of the probabilities of CRS and DLT_o , as illustrated in Scenarios 3 and 4. In Scenario 2, defined with a higher association between the CRS and the DLT_o , the DRtox_copula and DRtox_cond seem to perform better than the DRtox_indep as they account for the dependence between both toxicities. Indeed, in this scenario, more patients experience both toxicities (CRS and DLT_o) and therefore the DRtox_indep overestimates the probability of DLT. In Scenarios 1 to 4, the DRtox_copula has a higher proportion of trials in which the overdosing regimen S_5 is recommended as the MTD-regimen than the DRtox_indep and DRtox_cond, but this proportion is lower than or similar to that of the CRM.

In Scenario 5, where the true MTD-regimen is S_6 , the CRM has higher PCS than those of the DRtox_indep and DRtox_cond but similar results to those of the DRtox_copula. In this scenario the MTD-regimen is the last dose regimen of the panel; therefore fewer DLT, particularly fewer DLT_o , are observed. Indeed, on average, 5 DLT with 1.9 DLT_o occur per trial in Scenario 5, vs 7.7 DLT and 3.8 DLT_o in Scenario 1. In this scenario, the higher performance of the DRtox_copula and the CRM can be explained by their tendency to recommend higher dose regimens which was observed in Scenarios 1 to 4.

Scenario 6 was built on Set B which included an accelerated part for the first dose regimens to rapidly reach high steady-state doses. In this scenario, where S_2 is the true MTD-regimen, all three joint models yield similar results and still outperform the CRM.

All three joint models still perform better than the CRM in cases of independence or negative association between the CRS and the DLT_o (Web Appendix A.3 of Data S1).

We evaluated the influence of the prior distribution by comparing the results obtained with a prior ESS of 0.2, 2, and 7, the results can be found in Web Appendix B.1 of Data S1. Increasing the prior ESS leads to better results when the prior guesses of DLT probabilities are close to the truth (Scenarios 1-4 where S_4 is the true MTD-regimen), but also when the initial guesses of DLT probabilities underestimate the true DLT probabilities (Scenario 6 where S_2 is the true MTD-regimen). However, increasing the prior ESS leads to poorer results when the initial guesses of DLT probabilities overestimate the true DLT probabilities (Scenario 5 where S_6 is the true MTD-regimen). We also evaluated the impact of considering a gamma distribution for the DLT_o model and a log-normal distribution for the CRS model, and observed only a limited impact (Web Appendix B.2 of Data S1). We also evaluated the effect of a noninformative distribution on the DRtox_copula parameter, that is, $\gamma(0.1, 0.1)$. In this case, the results of the DRtox_copula are similar to the DRtox_indep (Web Appendix B.3 of Data S1). We also evaluated our methods after an empiric CRM and observed only limited impact on our proposed methods (Web Appendix B.4 of Data S1). Finally, we studied our methods on another set of dose regimens that is very different from the motivating trial to vary the timing of toxicity occurrences (Web Appendix C of Data S1). In this new set, our three joint models and the CRM have similar PCS so our three joint models do not improve the PCS but they can predict the probabilities of DLT of new dose regimens, and this will be developed in Section 3.2.3.

3.2.2 | Estimation of the toxicity curves

Gerard et al³³ showed that when only CRS occur, the DRtox leads to a better estimation of the entire dose regimen CRS relationship than the CRM. When enabling the occurrence of DLT_o , the probability of CRS is still well estimated with the DRtox. The joint probability of DLT is also well estimated around the MTD-regimen for all three joint approaches and

TABLE 2 Proportions of selecting each dose regimen S_k as the MTD-regimen over the 1000 trials in the six main toxicity scenarios. For each scenario, the marginal probabilities of dose limiting toxicity (DLT), cytokine release syndrome (CRS), and other DLT (DLT_0) are defined, and the association between the CRS and DLT_0 is represented by the average risk ratio (RR). Results are presented for the three joint approaches (DRtox_indep, DRtox_copula, and DRtox_cond) and the continual reassessment method (CRM). The proportions of correct selection (PCS) of the MTD-regimen are represented in bold

Scenario	Set	RR	Method	S_1	S_2	S_3	S_4	S_5	S_6
1	A	1.85	p_T	0.10	0.14	0.18	0.30	0.45	0.60
			$p_T^{(1)}$	0.05	0.07	0.1	0.18	0.22	0.30
			$p_T^{(2)}$	0.06	0.08	0.1	0.17	0.34	0.50
			DRtox_indep	0	4	24	55	16	2
			DRtox_copula	0	2	20	55	20	3
			DRtox_cond	0	3	25	55	16	2
			Logistic CRM	0	4	22	46	22	5
2	A	5.91	p_T	0.10	0.13	0.18	0.30	0.42	0.53
			$p_T^{(1)}$	0.07	0.10	0.13	0.22	0.28	0.36
			$p_T^{(2)}$	0.04	0.07	0.10	0.21	0.39	0.52
			DRtox_indep	0	4	27	48	16	3
			DRtox_copula	0	3	17	52	24	5
			DRtox_cond	0	3	23	52	18	4
			Logistic CRM	0	4	18	42	26	10
3	A	1.81	p_T	0.11	0.15	0.18	0.30	0.45	0.59
			$p_T^{(1)}$	0.03	0.05	0.07	0.13	0.16	0.23
			$p_T^{(2)}$	0.08	0.11	0.13	0.22	0.39	0.54
			DRtox_indep	1	3	22	57	15	2
			DRtox_copula	1	2	16	58	20	2
			DRtox_cond	1	3	22	58	15	2
			Logistic CRM	1	5	22	48	19	5
4	A	1.90	p_T	0.09	0.13	0.17	0.30	0.44	0.59
			$p_T^{(1)}$	0.07	0.10	0.13	0.23	0.29	0.37
			$p_T^{(2)}$	0.03	0.04	0.06	0.12	0.27	0.43
			DRtox_indep	0	3	24	56	14	2
			DRtox_copula	0	2	19	56	20	4
			DRtox_cond	0	3	24	55	15	2
			Logistic CRM	0	4	20	46	23	7
5	A	1.97	p_T	0.03	0.04	0.05	0.11	0.17	0.30
			$p_T^{(1)}$	0.02	0.03	0.05	0.09	0.11	0.17
			$p_T^{(2)}$	0.00	0.01	0.01	0.02	0.08	0.18
			DRtox_indep	0	0	0	4	29	67
			DRtox_copula	0	0	0	3	20	77
			DRtox_cond	0	0	0	4	29	67
			Logistic CRM	0	0	0	2	21	77

(Continues)

TABLE 2 (Continued)

Scenario	Set	RR	Method	S_1	S_2	S_3	S_4	S_5	S_6
6	B	1.70	p_T	0.16	0.30	0.43	0.55	0.73	0.86
			$p_T^{(1)}$	0.09	0.17	0.25	0.33	0.38	0.48
			$p_T^{(2)}$	0.08	0.18	0.29	0.41	0.67	0.84
			DRtox_indep	20	64	15	1	0	0
			DRtox_copula	14	63	21	2	0	0
			DRtox_cond	19	65	15	1	0	0
			Logistic CRM	17	56	24	3	0	0

the CRM. Both the DRtox_indep and DRtox_copula estimate the marginal probability of DLT_0 , while the DRtox_cond estimates the conditional probability of DLT_0 given that no CRS has occurred. The latter conditional probability is well estimated by the DRtox_cond, while the marginal probability is underestimated by the DRtox_indep and DRtox_copula (and estimated to be similar to the previous conditional probability) in all six main scenarios. Indeed, estimating the marginal probability of DLT_0 requires estimating the conditional probability of DLT_0 given CRS. The latter probability can only be estimated when both toxicities occur at the same time as drug administration is stopped when a DLT (either a CRS or DLT_0) occurs. However, CRS tend to occur at the beginning of the dose regimen, mainly at the first and fourth administrations, while most DLT_0 occur starting from the fourth administration. Thus, both toxicities rarely occur at the same time, and CRS occurrence may prevent future DLT_0 from occurring. This topic will be further developed in the discussion. Figures that illustrate the estimation of the different toxicity probabilities can be found in the supporting information (Web Appendix A.2 of Data S1).

3.2.3 | Prediction of new dose regimens

Our three proposed joint approaches and the CRM are able to recommend the MTD-regimen in the panel of the trial. However, alternative dose regimens may increase the efficacy of the drug while maintaining an acceptable probability of toxicity. For example, dose regimens where a lower steady-state dose is administered sooner or where a higher steady-state dose is administered later could be of interest. As the DRtox_indep, DRtox_copula, and DRtox_cond model the relationship between the entire dose regimen and the probability of DLT, they can predict the probability of DLT of new dose regimens that were not administered in the trial. For example, in Scenario 1, $S_4 = (1, 5, 10, 45, 45, 45, 45) \mu\text{g/kg}$ is selected for the MTD-regimen in more than 50% of trials where the steady-state dose of $45 \mu\text{g/kg}$ is administered starting from the fourth administration. We are able to predict the probabilities of two new dose regimens. For dose regimen $S_{\text{new1}} = (5, 10, 30, 30, 30, 30, 30) \mu\text{g/kg}$, having a true DLT probability of 0.27, the steady-state dose is decreased to $30 \mu\text{g/kg}$ but is administered earlier, from t_3 . For dose regimen $S_{\text{new2}} = (1, 5, 10, 30, 60, 60, 60) \mu\text{g/kg}$, having a true DLT probability of 0.31, the steady-state dose is increased to $60 \mu\text{g/kg}$ but is administered later, from t_5 . The estimated probabilities of DLT of the six initial dose regimens of the trial and the predicted probabilities of DLT of S_{new1} and S_{new2} are displayed in Figure 3. All approaches yield good estimations of the DLT probability of S_4 , but only the DRtox_indep, DRtox_copula, and DRtox_cond can predict the DLT probabilities of S_{new1} and S_{new2} that are close to the target. The three joint approaches can therefore help the clinical team find alternative dose regimens that can be investigated for efficacy in further stages of the trial, for example, in an expansion cohort.

The DRtox_indep, DRtox_copula, and DRtox_cond have similar precision in predicting new dose regimens with DLT probabilities close to the target in case of various associations between the CRS and the DLT_0 (Web Appendix A.3 of Data S1).

3.3 | Example from a single simulated trial

We illustrate our proposed methods on one hypothetical trial inspired by the motivating example. The trial was simulated under Scenario 2 with the settings specified in Section 3.1: it was conducted using a modified logistic CRM where




FIGURE 3 Violin plots of the estimated probabilities of dose limiting toxicity (DLT) in Scenario 1 for the six dose regimens of the panel and two additional dose regimens (S_{new1} and S_{new2}), on 1000 trials with the three proposed joint approaches and the continual reassessment method (CRM). The predicted DLT probabilities of the new dose regimens are framed in dotted line. Horizontal lines on the density estimates represent the median and first and third quartiles of the distributions, and the plus sign represents the mean. The dashed line represents the toxicity target, and the solid line represents the true DLT probabilities




FIGURE 4 Inclusion process of the simulated trial under a modified continual reassessment method (CRM). The type of dose limiting toxicity (DLT) is specified by the type of point, and the administration of occurrence of each type of DLT is given. The global number of DLT observed in each cohort of three patients is provided under each horizontal bar

patients were included by cohorts of size 3 until a maximum of 30 patients included. The toxicity target was 0.30. The inclusion process and the timing of occurrence of each type of toxicity are represented in Figure 4. During the trial, 8 DLT were observed, with 4 CRS and 6 DLT_o. For example, at cohort 5, three patients were included at dose regimen $S_4 = (1, 5, 10, 50, 50, 50, 50)$ $\mu\text{g}/\text{kg}$ administered at days ($t_1 = 1, t_2 = 5, t_3 = 9, t_4 = 13, t_5 = 17, t_6 = 21, t_7 = 25$). One of the three patient experienced a DLT, which was both a CRS and a DLT_o at administration t_4 . As drug administration is stopped when a DLT occurs, this patient only received the dose regimen composed of the four initial doses: (1, 5, 10, 50) $\mu\text{g}/\text{kg}$. The CRM estimated (0.11, 0.15, 0.22, **0.30**, 0.39, 0.47) for the posterior mean probabilities of DLT of S_1 to S_6 , therefore S_4 was estimated as the MTD-regimen. We can apply our proposed methods at the end of the trial.

CRS model

1. The PK/PD models (Appendix B.1) are fitted to the sampled individual concentrations and cytokine responses. Population and individual PK/PD parameters as well as the individual maximum peak of cytokine are estimated (see Web Appendix A.1 of Data S1 for the fitted cytokine response on one patient).
2. The CRS model is then fitted, the estimated posterior means of $\beta_{0,1}$ and $\beta_{1,1}$ are -2.40 and 2.78 , respectively.
3. The peak of cytokine at each dose regimen is simulated for 600 hypothetical patients from the population PK/PD parameters. The peak of cytokine can also be predicted for dose regimens that were not tested in the trial.
4. The posterior probabilities of CRS are estimated. The marginal probabilities of CRS are estimated by the mean of the posterior distributions: (0.05, 0.07, 0.09, **0.17**, 0.22, 0.28) for S_1 to S_6 .

DLT_o model

1. The cumulative marginal (and conditional given no CRS) probabilities of DLT_o are modeled with the cumulative dose received. The estimated posterior means of $\beta_{0,2}$ ($\beta_{0,2*}$) and $\beta_{1,2}$ ($\beta_{1,2*}$) are -1.39 (-1.88) and 0.05 (0.31), respectively.
2. The marginal posterior probabilities of DLT_o (and conditional posterior probabilities of DLT_o given no CRS) are estimated. The marginal probabilities of DLT_o (and the conditional probabilities of DLT_o given no CRS) are estimated by the mean of the posterior distributions: (0.11 (0.06), 0.13 (0.08), 0.15 (0.09), **0.21 (0.15)**, 0.33 (0.28), 0.41 (0.39)) for S_1 to S_6 .

DLT model

1. **DRtox_indep**: The posterior DLT probabilities are estimated from the marginal posterior probabilities of CRS and DLT_o. The marginal probabilities of DLT are estimated by the mean of the posterior distributions: (0.15, 0.19, 0.23, **0.34**, 0.48, 0.58) for S_1 to S_6 .
2. **DRtox_copula**: The DLT probabilities are modeled from the mean estimates of the marginal probabilities of CRS and DLT_o with the Clayton copula. The marginal probabilities of DLT are estimated by the mean of the posterior distributions: (0.13, 0.17, 0.20, **0.30**, 0.43, 0.52) for S_1 to S_6 .
3. **DRtox_cond**: The posterior DLT probabilities are estimated from the marginal posterior probabilities of CRS and the conditional posterior probabilities of DLT_o given no CRS. The marginal probabilities of DLT are estimated by the mean of the posterior distributions: (0.11, 0.14, 0.18, **0.29**, 0.44, 0.56) for S_1 to S_6 .

A figure of the posterior distributions can be found in Web Appendix A.1 of Data S1. The CRM and our three approaches lead to the same MTD-regimen recommendation, that is S_4 . As two patients experienced both a CRS and a DLT_o, the DRtox_indep overestimates the probability of DLT of S_4 while the DRtox_copula and DRtox_cond correctly estimate it because they account for the correlation between both toxicities. Moreover, as our methods model the relationship between the entire dose regimen and the probability of DLT, they can predict the probability of DLT of additional dose regimens that were not administered in the trial. For example, we can define the following additional dose regimen (5, 10, 30, 30, 30, 30, 30) $\mu\text{g}/\text{kg}$ than may have a clinical relevance because it reaches a lower value of the steady-state dose but earlier than S_4 . The DRtox_indep, DRtox_copula, and DRtox_cond predict 0.29, 0.26, and 0.24 for the probability of DLT of this new dose regimen, respectively. Therefore, at the end of the trial, our proposed methods support the recommendation of the CRM, and are able to propose an alternative dose regimen that could be evaluated in an expansion cohort.

4 | DISCUSSION

The aim of this work was to model the probability of DLT defined from the occurrence of two types of toxicity (CRS and DLT_o) in multiple dose administrations, that is, when dose regimens were administered, assuming that the CRS and the DLT_o might be associated. We developed a modeling approach for the DLT_o in addition to the existing DRtox approach that provides the CRS modeling with PK/PD. We proposed three approaches to model the joint distribution of DLT using the

independence assumption (DRtox_indep), a copula model (DRtox_copula), and a conditional formulation (DRtox_cond). Through an extensive simulation study, the three joint approaches were able to recommend the MTD-regimen in a set of dose regimens and to predict the DLT probability of new dose regimens that could be investigated in later stages of the trial. These methods were developed to be used when the dose allocation process is finished and once all data are gathered, especially once PK/PD measurements have already been analyzed using a population PK/PD approach. However, our methods should not be applied if the trial is stopped very early in the inclusion process and when not enough data is gathered to guaranty a feasible and/or reliable estimation of the parameters.

In most cases, the PCS of the three proposed methods were higher than that of the CRM. When the interaction between the DLT_0 and CRS was increased, the DRtox_copula and DRtox_cond gave better results as these two approaches account for the association between toxicities. When the MTD-regimen was the last dose regimen of the set and therefore few DLT were expected, the CRM had better performance than our proposed methods. In this case, as very few DLT_0 were observed, distinguishing between the different dose regimens was challenging for the cumulative DLT_0 model built on the cumulative dose. A simpler model that does not account for repeated doses could be fit on the DLT_0 to increase the PCS, but with such a model, prediction of new dose regimens would no longer be possible (Web Appendix A.4 of Data S1).

In addition to the potential gain in PCS compared to standard dose-escalation methods, one important feature of our proposed methods is that they were able to propose alternative dose regimens that could have a better clinical benefit than the existing dose regimens, in terms of efficacy for example. Indeed, our joint approaches model the relationship between the entire dose regimen and the probability of DLT and can therefore predict the DLT probability of new dose regimens. This prediction was accurate in our simulation settings, but is dependent on the number of patients included and the number of DLTs observed, as well as on the confidence in the PK/PD estimates. These potential new dose regimens need to be discussed by the trial stakeholders to determine the most suitable ones to be included in the expansion cohorts.

One limitation of the DRtox_indep and DRtox_copula was that the marginal probability of DLT_0 seemed to be underestimated (Web Appendix A.2 of Data S1). Indeed, from our motivating trial, CRS were expected to be observed at the beginning of the dose regimens, while DLT_0 occurred at the end. Once a CRS was observed, the administration of the drug was stopped, so there was no possibility left to observe any DLT_0 that might have occurred later in the planned drug administrations. This issue is different from censoring or competitive risks because the DLT_0 could be observed at the same time as the CRS. However, as drug administration was stopped, the patient did not receive the remaining administrations of the entire dose regimen that may have caused a DLT_0 . The DRtox_cond accounted for this issue as it modeled the conditional probability of DLT_0 given that no CRS had occurred. A second limitation was associated with the use of copula models. Cunanan and Koopmeiners⁵¹ showed that in the context of phase I/II trials, using a copula model instead of an independent model does not improve the performance in terms of correct dose selection, even in cases with highly correlated outcomes. They suggested adding a strong prior information on the correlation parameter of the copula models. In our case, we used an informative prior ($\gamma(1, 1)$) and observed that when using a noninformative prior ($\gamma(0.1, 0.1)$), the performance of the model could be decreased. We also observed that the DRtox_copula was more likely to recommend overdosing regimens than the DRtox_indep and DRtox_cond. Finally, as our modeling strategy requires multiple steps of modeling, it should only be considered when enough data is gathered. It should also be performed in close collaboration with pharmacometricians to be more confident about PK/PD models and estimates.

In conclusion, in this work we proposed to model a bivariate binary toxicity outcome in the context of an inpatient dose-escalation early phase trial in oncology. We accounted for all types of DLT while allowing the explicit modeling of the main type of toxicity expected (CRS) with PK/PD. Some dose-finding methods have been proposed that include PK measures.^{26,52} However, as PK/PD measures are analyzed by batch (and therefore are not available in real time) and the PK/PD models are developed after several cohorts of patients, a full sequential dose-allocation approach can be difficult to perform in practice. Our approaches were developed to be computed at the end of the trial, so using all trial information, including toxicities and PK/PD measures, is most practical. We have shown that even if the early phase trial is designed with a standard dose-allocation method, modeling PK/PD and multiple dose administrations at the end of the trial can allow the study stakeholders to make decisions about future development using all available information within a single modeling process.

ACKNOWLEDGEMENTS

This work was partially funded by a grant from the Association Nationale de la Recherche et de la Technologie, with Sanofi-Aventis R&D, Convention industrielle de formation par la recherche number 2018/0530. The authors would like to thank the three referees for their constructive comments that led to significant improvements of the manuscript.

CONFLICT OF INTEREST

The authors declare no potential conflict of interests.

ORCID

Emma Gerard  <https://orcid.org/0000-0003-2498-4801>

Sarah Zohar  <https://orcid.org/0000-0002-8429-2340>

Moreno Ursino  <https://orcid.org/0000-0002-5709-4322>

REFERENCES

1. Common terminology criteria for adverse events (CTCAE); 2021. https://ctep.cancer.gov/protocoldevelopment/electronic_applications/ctc.htm#ctc_60. Accessed April 14, 2021.
2. Storer BE. Design and analysis of phase I clinical trials. *Biometrics*. 1989;45(3):925-937.
3. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics*. 1990;46(1):33-48.
4. Topp MS, Gokbuget N, Zugmaier G, et al. Phase II trial of the anti-CD19 bispecific T cell-engager blinatumomab shows hematologic and molecular remissions in patients with relapsed or refractory B-precursor acute lymphoblastic leukemia. *J Clin Oncol*. 2014;32(36):4134-4140.
5. Topp MS, Gökbuget N, Stein AS, et al. Safety and activity of blinatumomab for adult patients with relapsed or refractory B-precursor acute lymphoblastic leukaemia: a multicentre, single-arm, phase 2 study. *Lancet Oncol*. 2015;16(1):57-66.
6. Uy GL, Aldoss I, Foster MC, et al. Flotetuzumab as salvage immunotherapy for refractory acute myeloid leukemia. *Blood*. 2021;137(6):751-762.
7. Preliminary results from a phase I study of APVO436; 2021. <https://ash.confex.com/ash/2020/webprogram/Paper141619.html>. Accessed April 16, 2021.
8. Complete responses in relapsed/refractory acute myeloid leukemia (AML) patients on a weekly dosing schedule of vibecotamab (XmAb14045); 2021. <https://ash.confex.com/ash/2020/webprogram/Paper134746.html>. Accessed April 16, 2021.
9. Regimen Definition National Cancer Institute website; 2021. <https://www.cancer.gov/publications/dictionaries/cancer-terms/def/regimen>. Accessed April 14, 2021.
10. Bekele BN, Thall PF. Dose-finding based on multiple toxicities in a soft tissue sarcoma trial. *J Am Stat Assoc*. 2004;99(465):26-35.
11. Lee S, Hershman D, Martin P, Leonard J, Cheung K. Validation of toxicity burden score for use in phase I clinical trials. *J Clin Oncol*. 2009;27(Suppl 15):2514-2514.
12. Ezzalfani M, Zohar S, Qin R, Mandrekar SJ, Deley MCL. Dose-finding designs using a novel quasi-continuous endpoint for multiple toxicities. *Stat Med*. 2013;32(16):2728-2746.
13. Lee SM, Cheng B, Cheung YK. Continual reassessment method with multiple toxicity constraints. *Biostatistics*. 2011;12(2):386-398.
14. Lin R. Bayesian optimal interval design with multiple toxicity constraints. *Biometrics*. 2018;74(4):1320-1330.
15. Liu S, Yuan Y. Bayesian optimal interval designs for phase I clinical trials. *J R Stat Soc Ser C Appl Stat*. 2015;64(3):507-523.
16. Lee SM, Ursino M, Cheung YK, Zohar S. Dose-finding designs for cumulative toxicities using multiple constraints. *Biostatistics*. 2019;20(1):17-29.
17. Thall PF, Nguyen HQ, Braun TM, Qazilbash MH. Using joint utilities of the times to response and toxicity to adaptively optimize schedule-dose regimens. *Biometrics*. 2013;69(3):673-682.
18. Wages NA, O'Quigley J, Conaway MR. Phase I design for completely or partially ordered treatment schedules. *Stat Med*. 2014;33(4):569-579.
19. Legedza AT, Ibrahim JG. Longitudinal design for phase I clinical trials using the continual reassessment method. *Control Clin Trials*. 2000;21(6):574-588.
20. Braun TM, Yuan Z, Thall PF. Determining a maximum-tolerated schedule of a cytotoxic agent. *Biometrics*. 2005;61(2):335-343.
21. Braun TM, Thall PF, Nguyen H, De Lima M. Simultaneously optimizing dose and schedule of a new cytotoxic agent. *Clin Trials*. 2007;4(2):113-124.
22. Liu CA, Braun TM. Parametric non-mixture cure models for schedule finding of therapeutic agents. *J R Stat Soc Ser C Appl Stat*. 2009;58(2):225-236.
23. Zhang J, Braun TM. A phase I Bayesian adaptive design to simultaneously optimize dose and schedule assignments both between and within patients. *J Am Stat Assoc*. 2013;108(503):892-901.
24. Fernandes LL, Taylor JM, Murray S. Adaptive phase I clinical trial design using Markov models for conditional probability of toxicity. *J Biopharm Stat*. 2016;26(3):475-498.
25. Ursino M, Biard L, Chevret S. Bayesian model for early dose-finding in phase I trials with multiple treatment courses; 2020. arXiv preprint arXiv:2012.03700.
26. Günhan BK, Weber S, Friede T. A Bayesian time-to-event pharmacokinetic model for phase I dose-escalation trials with multiple schedules. *Stat Med*. 2020;39(27):3986-4000.
27. SAR440234. National cancer institute website; 2020. <https://www.cancer.gov/publications/dictionaries/cancer-drug/def/anti-cd123-cd3-bite-antibody-sar440234?redirect=true>. Accessed October 10, 2020.

28. NCT03594955. ClinicalTrials.gov website; 2020. <https://clinicaltrials.gov/ct2/show/NCT03594955>. Accessed October 10, 2020.
29. Shimabukuro-Vornhagen A, Gödel P, Subklewe M, et al. Cytokine release syndrome. *J Immunother Cancer*. 2018;6(1):56.
30. Teachey DT, Lacey SF, Shaw PA, et al. Identification of predictive biomarkers for cytokine release syndrome after chimeric antigen receptor T-cell therapy for acute lymphoblastic leukemia. *Cancer Discov*. 2016;6(6):664-679.
31. Chen X, Kamperschroer C, Wong G, Xuan D. A modeling framework to characterize cytokine release upon T-cell-engaging bispecific antibody treatment: methodology and opportunities. *Clin Transl Sci*. 2019;12(6):600-608.
32. Boissel N, deBotton S, Thomas XG, et al. An open-label, first-in-human, dose escalation study of a novel CD3-CD123 bispecific T-cell engager administered as a single agent by intravenous infusion in patients with relapsed or refractory acute myeloid leukemia, B-cell acute lymphoblastic leukemia, or high risk myelodysplastic syndrome. *J Clin Oncol*. 2018;36(Suppl 15):TPS7076.
33. Gerard E, Zohar S, Thai HT, Lorenzato C, Riviere MK, Ursino M. Bayesian dose regimen assessment in early phase oncology incorporating pharmacokinetics and pharmacodynamics. *Biometrics*. 2021;1-13.
34. Wang Z, Han W. Biomarkers of cytokine release syndrome and neurotoxicity related to CAR-T cell therapy. *Biomark Res*. 2018;6(1):1-10.
35. Santomasso BD, Park JH, Salloum D, et al. Clinical and biological correlates of neurotoxicity associated with CAR T-cell therapy in patients with B-cell acute lymphoblastic leukemia. *Cancer Discov*. 2018;8(8):958-971.
36. Fried S, Avigdor A, Bielora B, et al. Early and late hematologic toxicity following CD19 CAR-T cells. *Bone Marrow Transplant*. 2019;54(10):1643-1650.
37. Siegler EL, Kenderian SS. Neurotoxicity and cytokine release syndrome after chimeric antigen receptor T cell therapy: insights into mechanisms and novel therapies. *Front Immunol*. 2020;11:1973.
38. Wei J, Liu Y, Wang C, et al. The model of cytokine release syndrome in CAR T-cell treatment for B-cell non-Hodgkin lymphoma. *Signal Transduct Target Ther*. 2020;5(1):1-9.
39. Nelsen RB. *An Introduction to Copulas*. Berlin, Germany: Springer Science & Business Media; 2007.
40. Braun TM. The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Control Clin Trials*. 2002;23(3):240-256.
41. Thall PF, Cook JD. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*. 2004;60(3):684-693.
42. Dragalin V, Fedorov V. Adaptive designs for dose-finding based on efficacy-toxicity response. *J Stat Plan Inference*. 2006;136(6):1800-1823.
43. Yin G, Yuan Y. A latent contingency table approach to dose finding for combinations of two agents. *Biometrics*. 2009;65(3):866-875.
44. Yin G, Yuan Y. Bayesian dose finding in oncology for drug combinations by copula regression. *J R Stat Soc Ser C Appl Stat*. 2009;58(2):211-224.
45. Perrone E, Müller WG. Optimal designs for copula models. *Statistics*. 2016;50(4):917-929.
46. Zhu M, Wu B, Brandl C, et al. Blinatumomab, a bispecific T-cell engager (BiTE®) for CD-19 targeted cancer immunotherapy: clinical pharmacology and its implications. *Clin Pharmacokinet*. 2016;55(10):1271-1288.
47. Morita S, Thall PF, Müller P. Determining the effective sample size of a parametric prior. *Biometrics*. 2008;64(2):595-602.
48. R Core Team. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing; 2018.
49. Lixoft *Monolix Version 2019 R1*. Antony, France: Lixoft SAS; 2019. <http://lixoft.com/products/monolix/>.
50. Stan Development Team *RStan: the R interface to Stan*; 2019. R package version 2.19.2.
51. Cunanan KM, Koopmeiners JS. Efficacy/toxicity dose-finding using hierarchical modeling for multiple populations. *Contemp Clin Trials*. 2018;71:162-172.
52. Ursino M, Zohar S, Lentz F, et al. Dose-finding methods for phase I clinical trials using pharmacokinetics in small populations. *Biom J*. 2017;59(4):804-825.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Gerard E, Zohar S, Lorenzato C, Ursino M, Riviere M-K. Bayesian modeling of a bivariate toxicity outcome for early phase oncology trials evaluating dose regimens. *Statistics in Medicine*. 2021;1-19. <https://doi.org/10.1002/sim.9113>

APPENDIX A. PRIOR ELICITATION

The elicitation of prior distributions is based on the initial guesses of the probabilities of DLT (π_1, \dots, π_K) of dose regimens (S_1, \dots, S_K). S_{k_T} is the dose regimen of reference, where $\pi_{k_T} = \delta_T$. We initially assume that the probabilities of CRS and DLT₀ are independent and equal, meaning that $\pi_k = 1 - \left(1 - \pi_k^{(1)}\right) \left(1 - \pi_k^{(2)}\right) = 1 - \left(1 - \pi_k^{(1)}\right)^2 = 1 - \left(1 - \pi_k^{(2)}\right)^2$, where $\pi_k^{(1)}$ and $\pi_k^{(2)}$ are the initial guesses of the probabilities of CRS and DLT₀, respectively. Therefore $\pi_k^{(1)} = \pi_k^{(2)} = 1 - \sqrt{1 - \pi_k}$.

For the CRS model, we consider $\beta_{0,1} \sim \mathcal{N}(\bar{\beta}_{0,1}, \sigma_{\beta_{0,1}}^2)$ and $\beta_{1,1} \sim \gamma(\alpha_1, \frac{\alpha_1}{\bar{\beta}_{1,1}})$, where $\bar{\beta}_{0,1} = \mathbb{E}[\beta_{0,1}]$ and $\bar{\beta}_{1,1} = \mathbb{E}[\beta_{1,1}]$. By construction, we have $\bar{\beta}_{0,1} = \text{logit} \left(1 - \sqrt{1 - \delta_T} \right)$ with $r_i^M = \bar{r}_{k_T}^M$. We can determine $\bar{\beta}_{1,1}$ using either only one dose regimen (differing from the reference dose regimen S_{k_T}), as follows:

$$\text{logit}^{-1} \left(\bar{\beta}_{0,1} + \bar{\beta}_{1,1} \log \left(\frac{\bar{r}_k^M}{\bar{r}_{k_T}^M} \right) \right) = \pi_k^{(1)}, \quad (\text{A1})$$

where $k \in \{1, \dots, K\}$ and $k \neq k_T$, or multiple dose regimens, such as the neighbors of the reference dose regimen, as follows:

$$\bar{\beta}_{1,1} = \arg \min_{\beta_{1,1}} \sum_{k=k_T-1}^{k_T+1} \left(\pi_k^{(1)} - \text{logit}^{-1} \left(\bar{\beta}_{0,1} + \beta_{1,1} \log \left(\frac{\bar{r}_k^M}{\bar{r}_{k_T}^M} \right) \right) \right)^2. \quad (\text{A2})$$

For the DLT₀ model, we consider $\beta_{0,2} \sim \mathcal{N}(\bar{\beta}_{0,2}, \sigma_{\beta_{0,2}}^2)$ and $\beta_{1,2} \sim \mathcal{N}(\bar{\beta}_{1,2}, \sigma_{\beta_{1,2}}^2)$, where $\bar{\beta}_{0,2} = \mathbb{E}[\beta_{0,2}]$ and $\bar{\beta}_{1,2} = \mathbb{E}[\beta_{1,2}]$. By construction, we have $\bar{\beta}_{0,2} = \text{logit} \left(1 - \sqrt{1 - \delta_T} \right)$ with $\sum_{l=1}^J d_{i,l} = D_{k_T}$. We can determine $\bar{\beta}_{1,2}$ using either only one dose regimen (differing from the reference dose regimen S_{k_T}), as follows:

$$\text{logit}^{-1} \left(\bar{\beta}_{0,2} + \exp(\bar{\beta}_{1,2}) \log \left(\frac{1}{D_{k_T}} \sum_{l=1}^J d_{k,l} \right) \right) = \pi_k^{(2)}, \quad (\text{A3})$$

where $k \in \{1, \dots, K\}$ and $k \neq k_T$, or multiple dose regimens, such as the neighbors of the reference dose regimen, as follows:

$$\bar{\beta}_{1,2} = \arg \min_{\beta_{1,2}} \sum_{k=k_T-1}^{k_T+1} \left(\pi_k^{(2)} - \text{logit}^{-1} \left(\bar{\beta}_{0,2} + \exp(\beta_{1,2}) \log \left(\frac{1}{D_{k_T}} \sum_{l=1}^J d_{k,l} \right) \right) \right)^2. \quad (\text{A4})$$

APPENDIX B. SIMULATION STUDY

B.1 PK/PD models

The PK/PD models were inspired by published models on blinatumomab. For the PK, we considered a 1-compartment infusion model,⁴⁶ where the drug concentration was defined as follows:

$$C(t) = \sum_{j=1}^J \mathbb{1}_{\{t-t_j > T_{infj}\}} \frac{d_{kj}}{T_{infj}} \frac{1}{kV} \left(1 - e^{-kT_{infj}} \right) e^{-k(t-t_j-T_{infj})} + \mathbb{1}_{\{t-t_j \leq T_{infj} \ \& \ t-t_j \geq 0\}} \frac{d_{kj}}{T_{infj}} \frac{1}{kV} \{ 1 - e^{-k(t-t_j)} \}, \quad (\text{B1})$$

where T_{infj} is the duration of the infusion of the j^{th} administration, V is the distribution volume, Cl is the clearance of elimination, and k is the micro-constant defined as $k = \frac{Cl}{V}$. We assumed that the delay between successive doses was greater than the infusion duration, meaning that $t_{j+1} - t_j > T_{infj}$ for $j \in \{1, \dots, J-1\}$.

For the PD, the objective was to model cytokine mitigation when inpatient dose-escalation was implemented. We simplified the model developed by Chen et al,³¹ which assumes that cytokine production is stimulated by the drug concentration but inhibited by cytokine exposure through the AUC. We defined the cytokine response as follows:

$$\frac{dE(t)}{dt} = \frac{E_{\max} C(t)^H}{EC_{50}^H + C(t)^H} \left\{ 1 - \frac{I_{\max} \text{AUC}_E(t)}{\frac{IC_{50}}{K^{J-1}} + \text{AUC}_E(t)} \right\} - k_{\text{deg}} E(t) \quad (\text{B2})$$

where $\text{AUC}_E(t)$ is the cumulative cytokine exposure.

For both the PK and PD models, we considered a proportional error model with a value of 0.1. The values of the PK/PD parameters used for the simulations were inspired by the estimated parameters of blinatumomab^{31,46} and are displayed in

TABLE B1 Definition and value of the pharmacokinetics and pharmacodynamics (PK/PD) parameters used for the simulation study. Parameter estimate represents the fixed effect, and the coefficient of variation (CV) is the square root of the diagonal of the variance-covariance matrix. These values are inspired by the parameters estimated on blinatumomab,^{31,46} with a modification of I_{\max} to observe cytokine mitigation after several administrations

	Parameter	Estimate (% CV)	Unit	Description
PK model	Cl	1.36 (41.9)	L/h	Clearance of elimination
	V	3.4 (0)	L	Volume of distribution
PD model	E_{\max}	$3.59 \cdot 10^5$ (14)	pg/mL/h	Maximum cytokine release rate
	EC_{50}	1.10^4 (0)	ng/mL	Drug exposure for half-maximum release
	H	0.92 (3)		Hill coefficient for cytokine release
	I_{\max}	0.995 (0)		Maximum inhibition of cytokine release
	IC_{50}	$1.82 \cdot 10^4$ (12)	pg/mL·h	Cytokine exposure for half-maximum inhibition
	k_{deg}	0.18 (13)	h^{-1}	Degradation rate for cytokine
	K	2.83 (36)		Priming factor for cytokine release

Table B1. For the estimation, we considered the parameters EC_{50} , I_{\max} and IC_{50} fixed and no random effects on V and H . Estimation was performed with Monolix software (Lixoft SAS., 2019).

B.2 Computation of toxicity probabilities

The true marginal probability of DLT was computed as follows:

$$\begin{aligned}
 p_T(S_k) &= 1 - \mathbb{P}(Y^{(1)} = 0, Y^{(2)} = 0) \\
 &= 1 - \mathbb{P}(Y^{(1)} = 0; Y_1^{(2)} = 0, \dots, Y_J^{(2)} = 0) \\
 &= 1 - \mathbb{P}(Y_1^{(2)} = 0, \dots, Y_J^{(2)} = 0 | Y^{(1)} = 0) \mathbb{P}(Y^{(1)} = 0) \\
 &= 1 - \prod_{j=1}^J \mathbb{P}(Y_j^{(2)} = 0 | Y_{j-1}^{(2)} = 0, \dots, Y_1^{(2)} = 0; Y^{(1)} = 0) \mathbb{P}(Y^{(1)} = 0) \\
 &= 1 - \prod_{j=1}^J \left(1 - \mathbb{P}(Y_j^{(2)} = 1 | Y_{j-1}^{(2)} = 0, \dots, Y_1^{(2)} = 0; Y^{(1)} = 0)\right) (1 - \mathbb{P}(Y^{(1)} = 1)) \\
 &= 1 - \prod_{j=1}^J \left(1 - \frac{1}{\lambda} f_{\theta}(d_{k,j}, D_{k,j-1})\right) (1 - p_T^{(1)}(S_k)). \tag{B3}
 \end{aligned}$$

The true marginal probability of CRS was computed by simulations using a Monte-Carlo method. The marginal probability of DLT₀ was computed as follows:

$$\begin{aligned}
 p_T^{(2)}(S_k) &= 1 - \mathbb{P}(Y_1^{(2)} = 0, \dots, Y_J^{(2)} = 0) \\
 &= 1 - \mathbb{P}(Y_1^{(2)} = 0, \dots, Y_J^{(2)} = 0 | Y^{(1)} = 1) p_T^{(1)}(S_k) - \mathbb{P}(Y_1^{(2)} = 0, \dots, Y_J^{(2)} = 0 | Y^{(1)} = 0) (1 - p_T^{(1)}(S_k)) \\
 &= 1 - \prod_{j=1}^J \mathbb{P}(Y_j^{(2)} = 0 | (Y_{j-1}^{(2)} = 0, \dots, Y_1^{(2)} = 0); Y_i^{(1)} = 1) p_T^{(1)}(S_k) \\
 &\quad - \prod_{j=1}^J \mathbb{P}(Y_j^{(2)} = 0 | (Y_{j-1}^{(2)} = 0, \dots, Y_1^{(2)} = 0); Y_i^{(1)} = 0) (1 - p_T^{(1)}(S_k)) \\
 &= 1 - \prod_{j=1}^J (1 - f_{\theta}(d_{k,j}, D_{k,j-1})) p_T^{(1)}(S_k) - \prod_{j=1}^J \left(1 - \frac{1}{\lambda} f_{\theta}(d_{k,j}, D_{k,j-1})\right) (1 - p_T^{(1)}(S_k)). \tag{B4}
 \end{aligned}$$

Chapter 6

Conclusion and discussion

In this thesis, we developed a Bayesian modeling approach, to be applied at the end of the dose-escalation stage of a phase I trial in oncology, to determine the maximum tolerated dose regimen (MTD-regimen) accounting for a complex assumption on the probability of toxicity after multiple administrations of the drug that can be related to a pharmacokinetic/pharmacodynamic (PK/PD) endpoint. This work was motivated by a first-in-human dose-escalation trial, implemented with a standard design, where complex dose regimens, defined with a fixed intra-patient dose-escalation scheme, have been implemented to decrease the risk of cytokine release syndrome (CRS) due to the mitigation of the peak of cytokine. Although CRS was expected to be the main type of DLT, other potential DLT, named DLT_o , could also occur but initial assumptions on the effect of complex dose regimens on the probability of DLT_o could not be raised.

The first part of the thesis focused on modeling the relationship between the dose regimen and the PD-related toxicity (CRS in the context of the motivating trial related to the peak of cytokine) by incorporating complete PK/PD modeling. We developed a dose regimen assessment (DRtox) approach that models the relationship between the dose regimen and the PK/PD endpoint and the relationship between the PK/PD endpoint and the probability of toxicity (Bayesian logistic or hierarchical model) and estimates the posterior probability of toxicity using a Monte Carlo method to include the interpatient PK/PD variability.

The second part of the thesis aimed at including the modeling of DLT_o while preserving the DRtox approach on the PD related toxicity in order to recommend the MTD-regimen accounting for all types of DLT. As no initial assumption could be raised on the effect of multiple drug administrations on the probability of DLT_o , contrary to first type of toxicity, we modeled the cumulative probability of DLT_o with the cumulative dose administered. We defined a DLT from the occurrence of at least a PD-related toxicity or DLT_o and developed three joint approaches to model the probability of DLT under various assumptions on the correlation between the PD-related toxicity and DLT_o : an independent model (DRtox_indep), a copula model (DRtox_copula) and a conditional model (DRtox_cond).

During extensive simulation studies, in both cases (modeling only the PD-related toxicity or including the modeling of DLT_o), our modeling approaches showed good performance in selecting the correct MTD-regimen in comparison with the standard dose-finding design implemented. Moreover, as our approaches modeled the relationship between the entire dose regimen and the probability of toxicity, they were

able to predict the probability of toxicity of dose regimens that were not evaluated in the trial but could be clinically relevant in terms of future efficacy for example. Our proposed modeling approaches can therefore support the choice of the MTD-regimen for future stages of the development using additional information gathered in early phase trials and could be applied on future clinical trials that evaluate complex dose regimens. We can nevertheless identify some limitations and provide several perspectives to extend our work.

The first limit of our modeling approaches is that they should be applied at the end of the dose-escalation stage of the trial, that is, when all data is collected, due to the incorporation of PK/PD modeling in the DRtox. Indeed, PK/PD data are usually analyzed by batch and might therefore not be available after each cohort of treated patients. Moreover, the PK/PD models need to be developed and reliable, so multiple cohorts of patients are required in order to fit reliably the model. The global approach that includes the modeling of DLT_o also requires to have enough data collected as the method is composed of at least two steps of modeling (three for DRtox_copula) including non-linear mixed effect models for PK/PD.

In our simulation studies, we assumed that the initial panel of dose regimens evaluated was ordered in terms of the probabilities of toxicity due to the fact that, when increasing the level of the dose regimen, a more aggressive intra-patient dose-escalation scheme or a higher steady state dose was defined. This complete ordering justified the implementation of standard dose-escalation designs, such as the 3+3 or the CRM, that rely on the monotonicity assumption. Although our approaches do not require the ordering of the dose regimens to be guaranteed, the risk of the PD-related toxicity is assumed to increase with the PK/PD endpoint while the risk of DLT_o is assumed to increase with the cumulative dose.

In our simulation study, we observed that the performance of the DRtox was impacted by the dose-escalation design chosen to conduct the trial and we therefore recommend to use more complex dose-escalation designs, such as the CRM, instead of the algorithm based 3+3 design, to have a better distribution of patients on the various dose-levels. As a perspective, we could evaluate multiple existing dose-escalation designs in the context of intra-patient dose-escalation or propose a simplified version of our approaches, that does not consider all assumptions, to provide better recommendations to conduct future trials while keeping our proposed approaches to be applied at the end of the trial. Moreover, we could fine-tune the characteristics of the DRtox to find, for example, the necessary number of patients and PD-related toxicities in order to apply our methods during the dose-escalation stage, however these characteristics would most likely depend on the PK/PD models considered.

The DRtox approach requires strong assumptions on the PK/PD mechanism generating the PD-related toxicity. Both the logistic (logistic DRtox) and the hierarchical (hierarchical DRtox) approaches were developed with only one predictor, which is the peak of one cytokine in the context of the motivating trial. However, this PD-related toxicity could be related to multiple, possibly correlated, variables, for example multiple cytokines in the context of the trial (IL-6, INF- γ). Extensions to multi-parameters models could be considered, but, as early phase trials include a small number of patients, it may be difficult to include multiple variables in the

model. An initial step of variable selection could then be required to identify the best PK/PD endpoint. Another idea could be to define the predictor from multiple PK/PD variables.

In our example, we focused on one PD endpoint that is the peak of cytokine. However, in this thesis, we proposed a global modeling framework in order to incorporate PK/PD modeling. This means that our modeling approach could consider various PK/PD endpoints, for example the AUC, assuming that the PK/PD responses can be well estimated in the trial (based on the fact that PK/PD models rely on biological knowledge). Moreover, we developed this modeling approach to find an acceptable dose regimen for all patients, but our methods could be extended for personalized dosing from the estimated individual PK/PD profiles.

In our simulation studies, we considered the same PK/PD models for simulation and estimation (even if in this step some parameters were fixed) as we assumed that even misspecified PK/PD models should be able to predict correctly the PD endpoint of interest. We were inspired by the PK/PD models developed on blinatumomab [184, 38], but additional simulations should be performed when the PK/PD models on the motivating trial become available. Overall, our method should be performed in close collaboration with pharmacometricians and further work could evaluate the impact of PK/PD model misspecifications.

Both versions of the DRtoxic had different assumptions. The logistic DRtoxic modeled the total risk of toxicity with the global PD-endpoint (maximum peak of cytokine in the context of the motivating trial) while the hierarchical DRtoxic modeled the risk of toxicity after each administration with the PD-endpoint observed after the administration (local peak of cytokine), it therefore assumed that toxicity occurred without delay. The hierarchical model was therefore not flexible if toxicity did not occur at the maximum peak but at a lower one. However, this situation was against our modeling assumptions which should be verified when data start to be available, but the logistic model could then be preferred to avoid too strict assumptions. In our simulation study, we simulated the CRS after an administration if the local peak of cytokine exceeded a threshold including inter patients variability but an alternative procedure of simulation could also be considered, for example defining the probability of toxicity for each peak of cytokine. In practice, modeling assumptions should be verified when first data become available.

Finally, in our work, the PK/PD modeling part and the toxicity modeling part were performed in two stages. PK/PD was modeled with standard frequentist approaches while toxicity was modeled using a Bayesian approach that includes PK/PD estimates. Our modeling proposal could therefore be applied in practice by dividing the modeling tasks between pharmacometricians and biostatisticians. An extension of this work could be to perform the PK/PD modeling in the Bayesian paradigm, and performing both modeling stages in a single step could be investigated.

In the second part of the thesis, we considered a simple model for the DLT_o as it was only one type of the possible DLTs. To account for the multiple administrations of varying doses of the dose regimen, we chose to model the cumulative probability of DLT_o with the cumulative dose administered. However, this simple model did not account for the timing of the administrations and the per-administration dose. In particular, this model could not differentiate dose regimens having the same

cumulative dose, for example $\mathbf{S}_1 = (5, 10, 15) \mu\text{g}/\text{kg}$ and $\mathbf{S}_2 = (10, 10, 10) \mu\text{g}/\text{kg}$. An extension of this model could be to consider a PK endpoint as a predictor instead of the cumulative dose. This extension could then account for patient's sensitivity to the drug and could then illustrate the potential correlation between the CRS and DLT_o . However, in case the AUC is considered, the dose regimens \mathbf{S}_1 and \mathbf{S}_2 defined earlier would also lead to the same cumulative AUC (if linear kinetics is assumed), but other endpoints could differentiate the dose regimens, for example the C_{max} .

When evaluating the results of the joint approaches, we observed that the `DRtox_indep` and `DRtox_copula` underestimated the marginal probability of DLT_o , but with a correct estimation of the probabilities of CRS and DLT. We observed this issue in our simulations because we assumed that CRS tend to occur at the beginning of the dose regimens (during the increasing lead-in dose) while DLT_o occur mostly at the end. Therefore, if a CRS occurred, the patient could not receive the remaining planned doses that may cause a DLT_o . We therefore mostly observed the DLT_o conditionally on the fact that the patient did not experience a CRS, unless both toxicities occurred after the same drug administration. This finding motivated us to develop the `DRtox_cond` approach where we modeled the conditional probability of DLT_o . This issue was different from competitive risks as the CRS did not prevent DLT_o to occur but prevented the patient to receive the remaining administrations that may cause a DLT_o . The opposite could also occur if a DLT_o was observed first, but it was not accounted for by the `DRtox_cond` approach.

To account for the potential association between toxicities in the joint DLT modeling, we proposed a copula and a conditional model. An alternative modeling approach could be to model the joint distribution of DLT using shared random effects [19].

In all methods, we recommended the MTD-regimen based on the point estimates of the probabilities of toxicity. Our methods can be easily extended to consider the entire posterior distributions and recommend the MTD-regimen based on the probabilities of under dosing, target interval, overdosing and unacceptable toxicity, similarly to Neuenschwander et al.[116]. However, this interval approach cannot be performed for the `DRtox_copula` as the copula model is fitted on the posterior means of the probabilities of CRS and DLT_o .

In our simulation studies, we assumed no dropout meaning that patients only received part of the dose regimen in case of DLT even though our methods account for the actual dose regimen administered, even an incomplete one. In practice, patients who leave the study prematurely after few drug administrations without experiencing a DLT may be considered as non DLT-evaluable.

In our work, we considered longitudinal toxicity endpoints, but an alternative modeling proposal could be to model time-to-event outcomes. A first idea could be to extend the work of Zhang and Braun [179] who developed a cure rate model, where the cure parameter is dose-dependent. Their model could be considered for the DLT_o assuming that the times to DLT_o after each administration are independent. To account for the complex toxicity assumption on the PD-related toxicity, the cure rate parameter could be modeled with the PD endpoint, similarly to the `DRtox` approach.

Another possibility could be to extend the work of Günhan et al. [67], where

the hazard of DLT_o could be modeled with drug concentration meaning that the probability of DLT_o would be modeled with the AUC of drug concentration. A similar approach could be proposed for the PD-related toxicity where the hazard of toxicity could be modeled with the PD response. However, this would mean that the probability of the PD related outcome is related to the AUC of the PD response, while in the motivating trial, the probability of CRS is assumed to be related to the peak of cytokine.

Finally, time to DLT_o and time to the PD-related toxicity could be modeled with Cox models with time varying outcomes. For the DLT_o , it could be modeled with the dose at each administration and for the PD-related outcome, a time dependent variable could be defined from the PD response.

In a future work, we could evaluate these time to event approaches in the context of small sample size and compare them with the methods proposed in this thesis.

In this work, we focused on modeling toxicity in early phase trials in oncology, but our approaches could be extended to other phases of the clinical trials. For example, if a PD endpoint can be related to efficacy, our approaches could be extended to phase II trials. Our approaches could also be considered for pediatric trials while incorporating prior information from adult trials.

In this thesis, incorporating PK/PD allowed us to account for a complex toxicity assumption and to evaluate various dose regimens. Our modeling proposal should be performed in close collaboration between statisticians and pharmacometricians to be applied successfully. As a more global perspective, we wish to foster more collaboration between statisticians and pharmacometricians during all stages of drug development, instead of considering both analyzes as complementary [81]. Indeed, incorporating PK/PD data and models more regularly in standard analyzes could improve the efficiency of drug development by considering all available biological knowledge [80, 66, 107]. We also believe that statisticians and pharmacometricians could benefit from an increased interaction and learning from each other [133].

Bibliography

- [1] APVO436: 2020 ASH Annual Meeting. <https://ash.confex.com/ash/2020/webprogram/Paper141619.html>. Accessed: 2021-06-15.
- [2] Blinatumomab. <https://www.cancer.gov/search/results?swKeyword=blinatumomab>. Accessed: 2021-06-15.
- [3] FDA: Blincyto approval. <https://www.fda.gov/news-events/press-announcements/fda-expands-approval-blincyto-treatment-type-leukemia-patients-who-have-certa>. Accessed: 2021-06-15.
- [4] FDA: The drug development process. <https://www.fda.gov/patients/learn-about-drug-and-device-approvals/drug-development-process>. Accessed: 2021-03-11.
- [5] Lixoft: Monolix Documentation Version 2020. <https://monolix.lixoft.com/>. Accessed: 2021-04-29.
- [6] MD Anderson Cancer Center, Biostatistics software. <https://biostatistics.mdanderson.org/SoftwareDownload/SingleSoftware/Index/72>. Accessed: 2021-05-19.
- [7] Monolix version 2019R1. Antony, France: Lixoft SAS, 2019. <http://lixoft.com/products/monolix/>.
- [8] NCI CTCAE. https://ctep.cancer.gov/protocoldevelopment/electronic_applications/ctc.htm#ctc_50. Accessed: 2021-04-26.
- [9] NCI: Definition of adverse event. <https://www.cancer.gov/publications/dictionaries/cancer-terms/def/adverse-event>. Accessed: 2021-04-26.
- [10] NCI: definition of cytokine. <https://www.cancer.gov/publications/dictionaries/cancer-terms/def/cytokine>. Accessed: 2021-06-07.
- [11] NCI: Definition of regimen. <https://www.cancer.gov/publications/dictionaries/cancer-terms/def/regimen>. Accessed: 2021-04-26.
- [12] NCI: Definition of treatment cycle. <https://www.cancer.gov/publications/dictionaries/cancer-terms/def/treatment-cycle>. Accessed: 2021-04-26.
- [13] NCI: Definition of treatment schedule. <https://www.cancer.gov/publications/dictionaries/cancer-terms/def/treatment-schedule?redirect=true>. Accessed: 2021-04-26.

- [14] NCT03594955. <https://clinicaltrials.gov/ct2/show/NCT03594955>. Accessed: 2021-06-07.
- [15] SAR440234. <https://www.cancer.gov/publications/dictionaries/cancer-drug/def/anti-cd123-cd3-bite-antibody-sar440234?redirect=true>. Accessed: 2021-06-07.
- [16] Stan Development Team. Stan Modeling Language Users Guide and Reference Manual, Version 2.27. <https://mc-stan.org>.
- [17] The SAEM algorithm: a powerful stochastic algorithm for population pharmacology modeling. <https://www.math.u-bordeaux.fr/MAS10/DOC/PDF-PRES/Lavielle.pdf>. Accessed: 2021-07-16.
- [18] PharmaIntelligence: 2021 Clinical Development Success Rates. <https://pharmaintelligence.informa.com/resources/product-content/2021-clinical-development-success-rates>, 2021. Accessed: 2021-04-26.
- [19] ALTZERINAKOU, M.-A., AND PAOLETTI, X. An adaptive design for the identification of the optimal dose using joint modeling of continuous repeated biomarker measurements and time-to-toxicity in phase I/II clinical trials in oncology. *Statistical methods in medical research* 29, 2 (2020), 508–521.
- [20] ANDRILLON, A., CHEVRET, S., LEE, S. M., AND BIARD, L. Dose-finding design and benchmark for a right censored endpoint. *Journal of Biopharmaceutical Statistics* 30, 6 (2020), 948–963.
- [21] BABB, J., ROGATKO, A., AND ZACKS, S. Cancer phase I clinical trials: efficient dose escalation with overdose control. *Statistics in medicine* 17, 10 (1998), 1103–1120.
- [22] 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. *Journal of biopharmaceutical statistics* 19, 3 (2009), 469–484.
- [23] BEKELE, B. N., AND THALL, P. F. Dose-finding based on multiple toxicities in a soft tissue sarcoma trial. *Journal of the American Statistical Association* 99, 465 (2004), 26–35.
- [24] BERRY, S. M., CARLIN, B. P., LEE, J. J., AND MULLER, P. *Bayesian adaptive methods for clinical trials*. CRC press, 2010.
- [25] BERTRAND, J., AND MENTRÉ, F. Mathematical expressions of the pharmacokinetic and pharmacodynamic models implemented in the Monolix software. *Paris Diderot University: Paris Diderot University* (2008).
- [26] BETANCOURT, M. A conceptual introduction to Hamiltonian Monte Carlo. *arXiv preprint arXiv:1701.02434* (2017).
- [27] BOISSEL, N., DE BOTTON, S., THOMAS, X. G., RAO, E., BONNEVAUX, H., RUBIN-CARREZ, C., GUERIF, S., BEYS, E., GOSSELIN, A., BAUCHET,

- A.-L., ET AL. An open-label, first-in-human, dose escalation study of a novel CD3-CD123 bispecific T-cell engager administered as a single agent by intravenous infusion in patients with relapsed or refractory acute myeloid leukemia, B-cell acute lymphoblastic leukemia, or high risk myelodysplastic syndrome., 2018.
- [28] BONATE, P. L., AND HOWARD, D. R. *Pharmacokinetics in Drug Development: Problems and Challenges in Oncology, Volume 4*, vol. 4. Springer, 2016.
- [29] BRADSHAW, E. L., SPILKER, M. E., ZANG, R., BANSAL, L., HE, H., JONES, R. D., LE, K., PENNEY, M., SCHUCK, E., TOPP, B., ET AL. Applications of quantitative systems pharmacology in model-informed drug discovery: perspective on impact and opportunities. *CPT: pharmacometrics & systems pharmacology* 8, 11 (2019), 777–791.
- [30] BRAUN, T. M. The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes. *Controlled clinical trials* 23, 3 (2002), 240–256.
- [31] BRAUN, T. M., AND JIA, N. A generalized continual reassessment method for two-agent phase I trials. *Statistics in biopharmaceutical research* 5, 2 (2013), 105–115.
- [32] BRAUN, T. M., KANG, S., AND TAYLOR, J. M. A phase I/II trial design when response is unobserved in subjects with dose-limiting toxicity. *Statistical methods in medical research* 25, 2 (2016), 659–673.
- [33] BRAUN, T. M., LEVINE, J. E., AND FERRARA, J. L. Determining a maximum tolerated cumulative dose: dose reassignment within the TITE-CRM. *Controlled clinical trials* 24, 6 (2003), 669–681.
- [34] BRAUN, T. M., THALL, P. F., NGUYEN, H., AND DE LIMA, M. Simultaneously optimizing dose and schedule of a new cytotoxic agent. *Clinical Trials* 4, 2 (2007), 113–124.
- [35] BRAUN, T. M., YUAN, Z., AND THALL, P. F. Determining a maximum-tolerated schedule of a cytotoxic agent. *Biometrics* 61, 2 (2005), 335–343.
- [36] BRUNTON, L. L., KNOLLMANN, B. C., AND HILAL-DANDAN, R. *Goodman & Gilman's: The pharmacological basis of therapeutics, 13e*. McGraw-Hill Education LLC., 2018.
- [37] CAI, C., YUAN, Y., AND JI, Y. A Bayesian dose-finding design for oncology clinical trials of combinational biological agents. *Journal of the Royal Statistical Society. Series C, Applied statistics* 63, 1 (2014), 159.
- [38] CHEN, X., KAMPERSCHROER, C., WONG, G., AND XUAN, D. A Modeling Framework to Characterize Cytokine Release upon T-Cell-Engaging Bispecific Antibody Treatment: Methodology and Opportunities. *Clinical and translational science* 12, 6 (2019), 600–608.

- [39] CHEN, Z., KRAILO, M. D., AZEN, S. P., AND TIGHIOUART, M. A novel toxicity scoring system treating toxicity response as a quasi-continuous variable in Phase I clinical trials. *Contemporary clinical trials* 31, 5 (2010), 473–482.
- [40] CHEUNG, Y. K. *Dose finding by the continual reassessment method*. CRC Press, 2011.
- [41] CHEUNG, Y. K., AND CHAPPELL, R. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* 56, 4 (2000), 1177–1182.
- [42] CHEVRET, S. *Statistical methods for dose-finding experiments*. Wiley, 2006.
- [43] COLLINS, J. M., GRIESHABER, C. K., AND CHABNER, B. A. Pharmacologically guided phase I clinical trials based upon preclinical drug development. *JNCI: Journal of the National Cancer Institute* 82, 16 (1990), 1321–1326.
- [44] COLLINS, J. M., ZAHARKO, D. S., DEDRICK, R. L., AND CHABNER, B. A. Potential roles for preclinical pharmacology in phase I clinical trials. *Cancer Treat Rep* 70, 1 (1986), 73–80.
- [45] COMETS, E. Etude de la réponse aux médicaments par la modélisation des relations dose-concentration-effet, 2010.
- [46] COMETS, E., AND ZOHAR, S. A survey of the way pharmacokinetics are reported in published phase I clinical trials, with an emphasis on oncology. *Clinical pharmacokinetics* 48, 6 (2009), 387–395.
- [47] COMMITTEE FOR MEDICINAL PRODUCTS FOR HUMAN USE AND OTHERS. Guideline on reporting the results of population pharmacokinetic analyses. *Doc. ref. CHMP/EWP/185990/06, London* (2007).
- [48] CONAWAY, M. R., DUNBAR, S., AND PEDDADA, S. D. Designs for single-or multiple-agent phase I trials. *Biometrics* 60, 3 (2004), 661–669.
- [49] COOK, N., HANSEN, A. R., SIU, L. L., AND RAZAK, A. R. A. Early phase clinical trials to identify optimal dosing and safety. *Molecular oncology* 9, 5 (2015), 997–1007.
- [50] CROWLEY, J., AND HOERING, A. *Handbook of statistics in clinical oncology*. CRC Press, 2012.
- [51] CUNANAN, K., AND KOOPMEINERS, J. S. Evaluating the performance of copula models in phase I-II clinical trials under model misspecification. *BMC Medical Research Methodology* 14, 1 (2014), 1–11.
- [52] DELYON, B., LAVIELLE, M., MOULINES, E., ET AL. Convergence of a stochastic approximation version of the EM algorithm. *The Annals of Statistics* 27, 1 (1999), 94–128.
- [53] DERENDORF, H., AND MEIBOHM, B. Modeling of pharmacokinetic/pharmacodynamic (PK/PD) relationships: concepts and perspectives. *Pharmaceutical research* 16, 2 (1999), 176–185.

- [54] DOUSSAU, A., THIÉBAUT, R., AND PAOLETTI, X. Dose-finding design using mixed-effect proportional odds model for longitudinal graded toxicity data in phase I oncology clinical trials. *Statistics in medicine* 32, 30 (2013), 5430–5447.
- [55] DRAGALIN, V., AND FEDOROV, V. Adaptive designs for dose-finding based on efficacy–toxicity response. *Journal of Statistical Planning and Inference* 136, 6 (2006), 1800–1823.
- [56] EZZALFANI, M., BURZYKOWSKI, T., AND PAOLETTI, X. Joint modelling of a binary and a continuous outcome measured at two cycles to determine the optimal dose. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 68, 2 (2019), 369–384.
- [57] EZZALFANI, M., ZOHAR, S., QIN, R., MANDREKAR, S. J., AND DELEY, M.-C. L. Dose-finding designs using a novel quasi-continuous endpoint for multiple toxicities. *Statistics in medicine* 32, 16 (2013), 2728–2746.
- [58] FAN, S. K., VENOOK, A. P., AND LU, Y. Design issues in dose-finding phase I trials for combinations of two agents. *Journal of biopharmaceutical statistics* 19, 3 (2009), 509–523.
- [59] FERNANDES, L. L., TAYLOR, J. M., AND MURRAY, S. Adaptive phase I clinical trial design using Markov models for conditional probability of toxicity. *Journal of biopharmaceutical statistics* 26, 3 (2016), 475–498.
- [60] FRIED, S., AVIGDOR, A., BIELORAI, B., MEIR, A., BESSER, M. J., SCHACHTER, J., SHIMONI, A., NAGLER, A., TOREN, A., AND JACOBY, E. Early and late hematologic toxicity following CD19 CAR-T cells. *Bone marrow transplantation* 54, 10 (2019), 1643–1650.
- [61] GARRALDA, E., DIENSTMANN, R., AND TABERNERO, J. Pharmacokinetic/pharmacodynamic modeling for drug development in oncology. *American Society of Clinical Oncology Educational Book* 37 (2017), 210–215.
- [62] GARRETT-MAYER, E. The continual reassessment method for dose-finding studies: a tutorial. *Clinical trials* 3, 1 (2006), 57–71.
- [63] GELMAN, A., CARLIN, J. B., STERN, H. S., DUNSON, D. B., VEHTARI, A., AND RUBIN, D. B. *Bayesian data analysis*. CRC press, 2013.
- [64] GERARD, E., ZOHAR, S., THAI, H.-T., LORENZATO, C., RIVIERE, M.-K., AND URSINO, M. Bayesian dose regimen assessment in early phase oncology incorporating pharmacokinetics and pharmacodynamics. *Biometrics* (2021), 1–13.
- [65] GEZMU, M., AND FLOURNOY, N. Group up-and-down designs for dose-finding. *Journal of statistical planning and inference* 136, 6 (2006), 1749–1764.
- [66] GIBSON, E., BRETZ, F., LOOBY, M., AND BORNKAMP, B. Key aspects of modern, quantitative drug development. *Statistics in Biosciences* 10, 2 (2018), 283–296.

- [67] GÜNHAN, B. K., WEBER, S., AND FRIEDE, T. A Bayesian time-to-event pharmacokinetic model for phase I dose-escalation trials with multiple schedules. *Statistics in Medicine* 39, 27 (2020), 3986–4000.
- [68] GÜNHAN, B. K., WEBER, S., SEROUTOU, A., AND FRIEDE, T. Phase I dose-escalation oncology trials with sequential multiple schedules. *BMC Medical Research Methodology* 21, 1 (2021), 1–14.
- [69] GUO, B., LI, Y., AND YUAN, Y. A dose–schedule finding design for phase I–II clinical trials. *Journal of the Royal Statistical Society. Series C, Applied statistics* 65, 2 (2016), 259.
- [70] GUO, W., WANG, S.-J., YANG, S., LYNN, H., AND JI, Y. A Bayesian interval dose-finding design addressing Ockham’s razor: mTPI-2. *Contemporary clinical trials* 58 (2017), 23–33.
- [71] HANSEN, A. R., COOK, N., RICCI, M. S., RAZAK, A., LE TOURNEAU, C., MCKEEVER, K., ROSKOS, L., DIXIT, R., SIU, L. L., AND HINRICH, M. J. Choice of starting dose for biopharmaceuticals in first-in-human phase I cancer clinical trials. *The oncologist* 20, 6 (2015), 653.
- [72] HUANG, B., AND KUAN, P. F. Time-to-event continual reassessment method incorporating treatment cycle information with application to an oncology phase I trial. *Biometrical Journal* 56, 6 (2014), 933–946.
- [73] 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. *Clinical Trials* 5, 5 (2008), 465–477.
- [74] IVANOVA, A. A new dose-finding design for bivariate outcomes. *Biometrics* 59, 4 (2003), 1001–1007.
- [75] IVANOVA, A., FLOURNOY, N., AND CHUNG, Y. Cumulative cohort design for dose-finding. *Journal of Statistical Planning and Inference* 137, 7 (2007), 2316–2327.
- [76] IVANOVA, A., AND WANG, K. A non-parametric approach to the design and analysis of two-dimensional dose-finding trials. *Statistics in Medicine* 23, 12 (2004), 1861–1870.
- [77] JI, Y., LI, Y., AND NEBIYOU BEKELE, B. Dose-finding in phase I clinical trials based on toxicity probability intervals. *Clinical trials* 4, 3 (2007), 235–244.
- [78] JI, Y., LIU, P., LI, Y., AND NEBIYOU BEKELE, B. A modified toxicity probability interval method for dose-finding trials. *Clinical trials* 7, 6 (2010), 653–663.
- [79] KOOPMEINERS, J. S., AND MODIANO, J. A bayesian adaptive phase i–ii clinical trial for evaluating efficacy and toxicity with delayed outcomes. *Clinical Trials* 11, 1 (2014), 38–48.

- [80] KOWALSKI, K. G. My career as a pharmacometrician and commentary on the overlap between statistics and pharmacometrics in drug development. *Statistics in Biopharmaceutical Research* 7, 2 (2015), 148–159.
- [81] KOWALSKI, K. G. Integration of pharmacometric and statistical analyses using clinical trial simulations to enhance quantitative decision making in clinical drug development. *Statistics in Biopharmaceutical Research* 11, 1 (2019), 85–103.
- [82] KUHN, E., AND LAVIELLE, M. Maximum likelihood estimation in nonlinear mixed effects models. *Computational statistics & data analysis* 49, 4 (2005), 1020–1038.
- [83] 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.
- [84] LE TOURNEAU, C., LEE, J. J., AND SIU, L. L. Dose escalation methods in phase I cancer clinical trials. *JNCI: Journal of the National Cancer Institute* 101, 10 (2009), 708–720.
- [85] LEE, D. W., GARDNER, R., PORTER, D. L., LOUIS, C. U., AHMED, N., JENSEN, M., GRUPP, S. A., AND MACKALL, C. L. Current concepts in the diagnosis and management of cytokine release syndrome. *Blood* 124, 2 (2014), 188–195.
- [86] LEE, J., THALL, P. F., JI, Y., AND MÜLLER, P. Bayesian dose-finding in two treatment cycles based on the joint utility of efficacy and toxicity. *Journal of the American Statistical Association* 110, 510 (2015), 711–722.
- [87] LEE, S., HERSHMAN, D., MARTIN, P., LEONARD, J., AND CHEUNG, Y. Toxicity burden score: a novel approach to summarize multiple toxic effects. *Annals of oncology* 23, 2 (2012), 537–541.
- [88] LEE, S. M., CHENG, B., AND CHEUNG, Y. K. Continual reassessment method with multiple toxicity constraints. *Biostatistics* 12, 2 (2011), 386–398.
- [89] LEE, S. M., AND CHEUNG, Y. K. Model calibration in the continual reassessment method. *Clinical Trials* 6, 3 (2009), 227–238.
- [90] LEGEDZA, A. T., AND IBRAHIM, J. G. Longitudinal design for phase I clinical trials using the continual reassessment method. *Controlled clinical trials* 21, 6 (2000), 574–588.
- [91] LI, D. H., WHITMORE, J. B., GUO, W., AND JI, Y. Toxicity and efficacy probability interval design for phase I adoptive cell therapy dose-finding clinical trials. *Clinical Cancer Research* 23, 1 (2017), 13–20.
- [92] LI, Y., BEKELE, B. N., JI, Y., AND COOK, J. D. Dose–schedule finding in phase I/II clinical trials using a Bayesian isotonic transformation. *Statistics in medicine* 27, 24 (2008), 4895–4913.

- [93] LIN, R. Bayesian optimal interval design with multiple toxicity constraints. *Biometrics* 74, 4 (2018), 1320–1330.
- [94] LIN, R., THALL, P. F., AND YUAN, Y. An adaptive trial design to optimize dose-schedule regimes with delayed outcomes. *Biometrics* 76, 1 (2020), 304–315.
- [95] LIN, R., AND YIN, G. Bayesian optimal interval design for dose finding in drug-combination trials. *Statistical methods in medical research* 26, 5 (2017), 2155–2167.
- [96] LIN, R., AND YIN, G. STEIN: A simple toxicity and efficacy interval design for seamless phase I/II clinical trials. *Statistics in medicine* 36, 26 (2017), 4106–4120.
- [97] LIN, X., AND JI, Y. Probability intervals of toxicity and efficacy design for dose-finding clinical trials in oncology. *Statistical Methods in Medical Research* 30, 3 (2021), 843–856.
- [98] LINDSTROM, M. J., AND BATES, D. M. Nonlinear mixed effects models for repeated measures data. *Biometrics* (1990), 673–687.
- [99] LIU, C. A., AND BRAUN, T. M. Parametric non-mixture cure models for schedule finding of therapeutic agents. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 58, 2 (2009), 225–236.
- [100] LIU, S., CAI, C., AND NING, J. Up-and-down designs for phase I clinical trials. *Contemporary clinical trials* 36, 1 (2013), 218–227.
- [101] LIU, S., AND YUAN, Y. Bayesian optimal interval designs for phase I clinical trials. *Journal of the Royal Statistical Society: Series C: Applied Statistics* (2015), 507–523.
- [102] LYU, J., CURRAN, E., AND JI, Y. Bayesian adaptive design for finding the maximum tolerated sequence of doses in multicycle dose-finding clinical trials. *JCO Precision Oncology* 2 (2018), 1–19.
- [103] MANDER, A. P., AND SWEETING, M. J. A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Statistics in medicine* 34, 8 (2015), 1261–1276.
- [104] 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. *Statistics in medicine* 26, 11 (2007), 2317–2330.
- [105] MAO, X., AND CHEUNG, Y. K. Sequential designs for individualized dosing in phase I cancer clinical trials. *Contemporary clinical trials* 63 (2017), 51–58.
- [106] MARAYATI, R., QUINN, C. H., AND BEIERLE, E. A. Immunotherapy in pediatric solid tumors—A systematic review. *Cancers* 11, 12 (2019), 2022.

-
- [107] MENTRÉ, F., FRIBERG, L. E., DUFFULL, S., FRENCH, J., LAUFFENBURGER, D. A., LI, L., MAGER, D. E., SINHA, V., SOBIE, E., AND ZHAO, P. Pharmacometrics and systems pharmacology 2030. *Clinical Pharmacology & Therapeutics* 107, 1 (2020), 76–78.
- [108] METER, E. M. V., GARRETT-MAYER, E., AND BANDYOPADHYAY, D. Proportional odds model for dose-finding clinical trial designs with ordinal toxicity grading. *Statistics in medicine* 30, 17 (2011), 2070–2080.
- [109] MORITA, S., THALL, P. F., AND MÜLLER, P. Determining the effective sample size of a parametric prior. *Biometrics* 64, 2 (2008), 595–602.
- [110] MOULD, D. R., AND UPTON, R. Basic concepts in population modeling, simulation, and model-based drug development. *CPT: pharmacometrics & systems pharmacology* 1, 9 (2012), 1–14.
- [111] MUENZ, D. G., BRAUN, T. M., AND TAYLOR, J. M. Modeling adverse event counts in phase I clinical trials of a cytotoxic agent. *Clinical Trials* 15, 4 (2018), 386–397.
- [112] MURPHY, S. A. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 65, 2 (2003), 331–355.
- [113] MURTAUGH, P. A., AND FISHER, L. D. Bivariate binary models of efficacy and toxicity in dose-ranging trials. *Communications in Statistics-Theory and Methods* 19, 6 (1990), 2003–2020.
- [114] NEBIYOU BEKELE, B., AND SHEN, Y. A Bayesian approach to jointly modeling toxicity and biomarker expression in a phase I/II dose-finding trial. *Biometrics* 61, 2 (2005), 343–354.
- [115] NELSEN, R. B. *An introduction to copulas*. Springer Science & Business Media, 2007.
- [116] NEUENSCHWANDER, B., BRANSON, M., AND GSPONER, T. Critical aspects of the Bayesian approach to phase I cancer trials. *Statistics in medicine* 27, 13 (2008), 2420–2439.
- [117] O’QUIGLEY, J., HUGHES, M. D., AND FENTON, T. Dose-finding designs for HIV studies. *Biometrics* 57, 4 (2001), 1018–1029.
- [118] O’QUIGLEY, J., IASONOS, A., AND BORNKAMP, B. *Handbook of methods for designing, monitoring, and analyzing dose-finding trials*. CRC Press, 2017.
- [119] O’QUIGLEY, J., PEPE, M., AND FISHER, L. Continual reassessment method: a practical design for phase 1 clinical trials in cancer. *Biometrics* (1990), 33–48.
- [120] PAN, H., LIN, R., ZHOU, Y., AND YUAN, Y. Keyboard design for phase I drug-combination trials. *Contemporary clinical trials* 92 (2020), 105972.
- [121] PAOLETTI, X., DOUSSAU, A., EZZALFANI, M., RIZZO, E., AND THIÉBAUT, R. Dose finding with longitudinal data: simpler models, richer outcomes. *Statistics in medicine* 34, 22 (2015), 2983–2998.

- [122] PARK, K. A review of modeling approaches to predict drug response in clinical oncology. *Yonsei medical journal* 58, 1 (2017), 1–8.
- [123] PATTERSON, S., FRANCIS, S., IRESON, M., WEBBER, D., AND WHITEHEAD, J. A novel Bayesian decision procedure for early-phase dose-finding studies. *Journal of biopharmaceutical statistics* 9, 4 (1999), 583–597.
- [124] PERRONE, E., AND MÜLLER, W. G. Optimal designs for copula models. *Statistics* 50, 4 (2016), 917–929.
- [125] PIANTADOSI, S., AND LIU, G. Improved designs for dose escalation studies using pharmacokinetic measurements. *Statistics in medicine* 15, 15 (1996), 1605–1618.
- [126] QUINTANA, M., LI, D. H., ALBERTSON, T. M., AND CONNOR, J. T. A Bayesian adaptive phase 1 design to determine the optimal dose and schedule of an adoptive T-cell therapy in a mixed patient population. *Contemporary clinical trials* 48 (2016), 153–165.
- [127] R CORE TEAM. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2020.
- [128] RATAIN, M., PLUNKETT JR, W., ET AL. Principles of pharmacokinetics. *Holland-Frei cancer medicine, 6th ed.* Decker, Hamilton, BC, Canada (2003).
- [129] RIVIERE, M.-K., DUBOIS, F., AND ZOHAR, S. Competing designs for drug combination in phase I dose-finding clinical trials. *Statistics in medicine* 34, 1 (2015), 1–12.
- [130] RIVIERE, M.-K., YUAN, Y., DUBOIS, F., AND ZOHAR, S. A Bayesian dose-finding design for drug combination clinical trials based on the logistic model. *Pharmaceutical statistics* 13, 4 (2014), 247–257.
- [131] ROGATKO, A., SCHOENECK, D., JONAS, W., TIGHIOUART, M., KHURI, F. R., AND PORTER, A. Translation of innovative designs into phase I trials. *Journal of Clinical Oncology* 25, 31 (2007), 4982–4986.
- [132] SANTOMASSO, B. D., PARK, J. H., SALLOUM, D., RIVIERE, I., FLYNN, J., MEAD, E., HALTON, E., WANG, X., SENECHAL, B., PURDON, T., ET AL. Clinical and biological correlates of neurotoxicity associated with CAR T-cell therapy in patients with B-cell acute lymphoblastic leukemia. *Cancer discovery* 8, 8 (2018), 958–971.
- [133] SENN, S. Statisticians and pharmacokineticists: what they can still learn from each other. *Clinical Pharmacology & Therapeutics* 88, 3 (2010), 328–334.
- [134] SHARMA, A., EBLING, W. F., AND JUSKO, W. J. Precursor-dependent indirect pharmacodynamic response model for tolerance and rebound phenomena. *Journal of pharmaceutical sciences* 87, 12 (1998), 1577–1584.
- [135] SHARMA, A., AND JUSKO, W. J. Characteristics of indirect pharmacodynamic models and applications to clinical drug responses. *British journal of clinical pharmacology* 45, 3 (1998), 229–239.

- [136] SHEINER, L. B., ROSENBERG, B., AND MELMON, K. L. Modelling of individual pharmacokinetics for computer-aided drug dosage. *Computers and Biomedical Research* 5, 5 (1972), 441–459.
- [137] SHIMABUKURO-VORNHAGEN, A., GÖDEL, P., SUBKLEWE, M., STEMMLER, H. J., SCHLÖSSER, H. A., SCHLAAK, M., KOCHANEK, M., BÖLL, B., AND VON BERGWELT-BAILDON, M. S. Cytokine release syndrome. *Journal for immunotherapy of cancer* 6, 1 (2018), 1–14.
- [138] SIEGLER, E. L., AND KENDERIAN, S. S. Neurotoxicity and cytokine release syndrome after chimeric antigen receptor T cell therapy: Insights Into mechanisms and novel therapies. *Frontiers in Immunology* 11 (2020), 1973.
- [139] SIMON, R., RUBINSTEIN, L., ARBUCK, S. G., CHRISTIAN, M. C., FREIDLIN, B., AND COLLINS, J. Accelerated titration designs for phase I clinical trials in oncology. *Journal of the National Cancer Institute* 89, 15 (1997), 1138–1147.
- [140] SINGH, A., DEES, S., AND GREWAL, I. S. Overcoming the challenges associated with CD3+ T-cell redirection in cancer. *British Journal of Cancer* (2021), 1–12.
- [141] SORGER, P. K., ALLERHEILIGEN, S. R., ABERNETHY, D. R., ALTMAN, R. B., BROUWER, K. L., CALIFANO, A., D'ARGENIO, D. Z., IYENGAR, R., JUSKO, W. J., LALONDE, R., ET AL. Quantitative and systems pharmacology in the post-genomic era: new approaches to discovering drugs and understanding therapeutic mechanisms. In *An NIH white paper by the QSP workshop group* (2011), vol. 48, NIH Bethesda Bethesda, MD, pp. 1–47.
- [142] STAAB, A., ROOK, E., MALIEPAARD, M., AARONS, L., AND BENSON, C. Modeling and simulation in clinical pharmacology and dose finding. *CPT: pharmacometrics & systems pharmacology* 2, 2 (2013), 1–3.
- [143] STALLARD, N., AND TODD, S. Seamless phase II/III designs. *Statistical methods in medical research* 20, 6 (2011), 623–634.
- [144] STORER, B. E. Design and analysis of phase I clinical trials. *Biometrics* (1989), 925–937.
- [145] TAKEDA, K., KOMATSU, K., AND MORITA, S. Bayesian dose-finding phase I trial design incorporating pharmacokinetic assessment in the field of oncology. *Pharmaceutical statistics* 17, 6 (2018), 725–733.
- [146] TAKEDA, K., TAGURI, M., AND MORITA, S. BOIN-ET: Bayesian optimal interval design for dose finding based on both efficacy and toxicity outcomes. *Pharmaceutical statistics* 17, 4 (2018), 383–395.
- [147] TEACHEY, D. T., LACEY, S. F., SHAW, P. A., MELENHORST, J. J., MAUDE, S. L., FREY, N., PEQUIGNOT, E., GONZALEZ, V. E., CHEN, F., FINKLESTEIN, J., ET AL. Identification of predictive biomarkers for cytokine release syndrome after chimeric antigen receptor T-cell therapy for acute lymphoblastic leukemia. *Cancer discovery* 6, 6 (2016), 664–679.

- [148] THALL, P. F., AND COOK, J. D. Dose-finding based on efficacy–toxicity trade-offs. *Biometrics* 60, 3 (2004), 684–693.
- [149] 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 (2003), 487–496.
- [150] THALL, P. F., NGUYEN, H. Q., BRAUN, T. M., AND QAZILBASH, M. H. Using joint utilities of the times to response and toxicity to adaptively optimize schedule–dose regimes. *Biometrics* 69, 3 (2013), 673–682.
- [151] 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* (1998), 251–264.
- [152] TOPP, M. S., GÖKBUGET, N., STEIN, A. S., ZUGMAIER, G., O’BIEN, S., BARGOU, R. C., DOMBRET, H., FIELDING, A. K., HEFFNER, L., LARSON, R. A., ET AL. Safety and activity of blinatumomab for adult patients with relapsed or refractory B-precursor acute lymphoblastic leukaemia: a multicentre, single-arm, phase 2 study. *The Lancet Oncology* 16, 1 (2015), 57–66.
- [153] TOPP, M. S., GOKBUGET, N., ZUGMAIER, G., KLAPPERS, P., STELLJES, M., NEUMANN, S., VIARDOT, A., MARKS, R., DIEDRICH, H., FAUL, C., ET AL. Phase II trial of the anti-CD19 bispecific T cell-engager blinatumomab shows hematologic and molecular remissions in patients with relapsed or refractory B-precursor acute lymphoblastic leukemia. *J Clin Oncol* 32, 36 (2014), 4134–4140.
- [154] URSINO, M., BIARD, L., AND CHEVRET, S. Bayesian model for early dose-finding in phase I trials with multiple treatment courses. *Biometrical Journal* (2021), 1–12.
- [155] URSINO, M., ZOHAR, S., LENTZ, F., ALBERTI, C., FRIEDE, T., STALLARD, N., AND COMETS, E. Dose-finding methods for phase I clinical trials using pharmacokinetics in small populations. *Biometrical Journal* 59, 4 (2017), 804–825.
- [156] US FOOD AND DRUG ADMINISTRATION AND OTHERS. Population Pharmacokinetics Guidance for Industry Draft Guidance. *US Department of Health and Human Services Food and Drug Administration Center for Drug Evaluation and Research (CDER) Center for Biologics Evaluation and Research (CBER)* (2019).
- [157] UY, G. L., ALDOSS, I., FOSTER, M. C., SAYRE, P. H., WIEDUWILT, M. J., ADVANI, A. S., GODWIN, J. E., ARELLANO, M. L., SWEET, K. L., EMADI, A., ET AL. Flotetuzumab as salvage immunotherapy for refractory acute myeloid leukemia. *Blood, The Journal of the American Society of Hematology* 137, 6 (2021), 751–762.
- [158] WAGES, N. A., AND CONAWAY, M. R. Specifications of a continual reassessment method design for phase I trials of combined drugs. *Pharmaceutical statistics* 12, 4 (2013), 217–224.

- [159] WAGES, N. A., CONAWAY, M. R., AND O'QUIGLEY, J. Continual reassessment method for partial ordering. *Biometrics* 67, 4 (2011), 1555–1563.
- [160] WAGES, N. A., O'QUIGLEY, J., AND CONAWAY, M. R. Phase I design for completely or partially ordered treatment schedules. *Statistics in medicine* 33, 4 (2014), 569–579.
- [161] WANG, C., CHEN, T. T., AND TYAN, I. Designs for phase I cancer clinical trials with differentiation of graded toxicity. *Communications in Statistics-Theory and Methods* 29, 5-6 (2000), 975–987.
- [162] WANG, K., AND IVANOVA, A. Two-dimensional dose finding in discrete dose space. *Biometrics* 61, 1 (2005), 217–222.
- [163] WANG, Z., AND HAN, W. Biomarkers of cytokine release syndrome and neurotoxicity related to CAR-T cell therapy. *Biomarker research* 6, 1 (2018), 1–10.
- [164] WEI, G. C., AND TANNER, M. A. Calculating the content and boundary of the highest posterior density region via data augmentation. *Biometrika* 77, 3 (1990), 649–652.
- [165] WEI, J., LIU, Y., WANG, C., ZHANG, Y., TONG, C., DAI, G., WANG, W., RASKO, J. E., MELENHORST, J. J., QIAN, W., ET AL. The model of cytokine release syndrome in CAR T-cell treatment for B-cell non-Hodgkin lymphoma. *Signal Transduction and Targeted Therapy* 5, 1 (2020), 1–9.
- [166] WHITEHEAD, J., PATTERSON, S., WEBBER, D., FRANCIS, S., AND ZHOU, Y. Easy-to-implement Bayesian methods for dose-escalation studies in healthy volunteers. *Biostatistics* 2, 1 (2001), 47–61.
- [167] WHITEHEAD, J., ZHOU, Y., HAMPSON, L., LEDENT, E., AND PEREIRA, A. A Bayesian approach for dose-escalation in a phase I clinical trial incorporating pharmacodynamic endpoints. *Journal of Biopharmaceutical statistics* 17, 6 (2007), 1117–1129.
- [168] WHITEHEAD, J., ZHOU, Y., STEVENS, J., BLAKEY, G., PRICE, J., AND LEADBETTER, J. Bayesian decision procedures for dose-escalation based on evidence of undesirable events and therapeutic benefit. *Statistics in medicine* 25, 1 (2006), 37–53.
- [169] YAN, F., MANDREKAR, S. J., AND YUAN, Y. Keyboard: a novel Bayesian toxicity probability interval design for phase I clinical trials. *Clinical Cancer Research* 23, 15 (2017), 3994–4003.
- [170] YAP, C., BILLINGHAM, L. J., CHEUNG, Y. K., CRADDOCK, C., AND O'QUIGLEY, J. Dose transition pathways: the missing link between complex dose-finding designs and simple decision-making. *Clinical Cancer Research* 23, 24 (2017), 7440–7447.
- [171] 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 (2006), 777–787.

- [172] YIN, G., AND YUAN, Y. A latent contingency table approach to dose finding for combinations of two agents. *Biometrics* 65, 3 (2009), 866–875.
- [173] YIN, G., AND YUAN, Y. Bayesian dose finding in oncology for drug combinations by copula regression. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 58, 2 (2009), 211–224.
- [174] YUAN, Y., LEE, J. J., AND HILSENBECK, S. G. Model-assisted designs for early-phase clinical trials: simplicity meets superiority. *JCO Precision Oncology* 3 (2019), 1–12.
- [175] YUAN, Y., NGUYEN, H. Q., AND THALL, P. F. *Bayesian designs for phase I–II clinical trials*. CRC Press, 2016.
- [176] YUAN, Y., AND YIN, G. Sequential continual reassessment method for two-dimensional dose finding. *Statistics in medicine* 27, 27 (2008), 5664–5678.
- [177] YUAN, Y., AND YIN, G. Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 58, 5 (2009), 719–736.
- [178] YUAN, Z., CHAPPELL, R., AND BAILEY, H. The continual reassessment method for multiple toxicity grades: a Bayesian quasi-likelihood approach. *Biometrics* 63, 1 (2007), 173–179.
- [179] ZHANG, J., AND BRAUN, T. M. A phase I Bayesian adaptive design to simultaneously optimize dose and schedule assignments both between and within patients. *Journal of the American Statistical Association* 108, 503 (2013), 892–901.
- [180] ZHANG, L., AND YUAN, Y. A practical Bayesian design to identify the maximum tolerated dose contour for drug combination trials. *Statistics in medicine* 35, 27 (2016), 4924–4936.
- [181] ZHANG, W., SARGENT, D. J., AND MANDREKAR, S. An adaptive dose-finding design incorporating both toxicity and efficacy. *Statistics in medicine* 25, 14 (2006), 2365–2383.
- [182] ZHOU, H., MURRAY, T. A., PAN, H., AND YUAN, Y. Comparative review of novel model-assisted designs for phase I clinical trials. *Statistics in Medicine* 37, 14 (2018), 2208–2222.
- [183] ZHOU, Y., WHITEHEAD, J., BONVINI, E., AND STEVENS, J. W. Bayesian decision procedures for binary and continuous bivariate dose-escalation studies. *Pharmaceutical Statistics: The Journal of Applied Statistics in the Pharmaceutical Industry* 5, 2 (2006), 125–133.
- [184] ZHU, M., WU, B., BRANDL, C., JOHNSON, J., WOLF, A., CHOW, A., AND DOSHI, S. Blinatumomab, a bispecific T-cell engager (BiTE®) for CD-19 targeted cancer immunotherapy: clinical pharmacology and its implications. *Clinical pharmacokinetics* 55, 10 (2016), 1271–1288.

Abstract

Title: Bayesian toxicity modeling for dose regimen assessment in early phase oncology trials incorporating pharmacokinetics and pharmacodynamics

Abstract: Phase I trials in oncology aim to evaluate the toxicity of a new drug and allow to gather information on the mechanism of action of the molecule through pharmacokinetic and pharmacodynamic (PK/PD) studies. They aim to determine the maximum tolerated dose (MTD) defined as the highest dose that does not exceed a predefined probability of dose-limiting toxicity (DLT) [42, 118]. DLT is a binary toxicity defined from the occurrence of severe adverse events. Sequential dose-escalation designs have been developed to determine the MTD after the first cycle of therapy, as the 3+3 [144] or the continual reassessment method (CRM) [119]. Evaluating dose regimens, defined as the combination of the schedule of administration and the dose, can reduce the risk of toxicity while preserving future efficacy. Some designs have been developed to evaluate the schedule of administration in addition to the dose, but they do not account for a complex assumption on toxicity mechanism [150, 160, 179]. In a phase I trial (NCT03594955 [14]), complex dose regimens, defined by multiple administrations of increasing doses until a fixed dose, have been implemented to reduce the risk of the main toxicity expected, the cytokine release syndrome (CRS). A standard dose-escalation design was conducted but it does not account for the information on the multiple administrations that compose the dose regimens. The aim of this thesis is to develop an approach, to be applied at the end of the trial, in order to determine the maximum tolerated dose regimen (MTD-regimen) in the context of the motivating trial.

The first part of the thesis is based on the motivating trial, where the reduction of the risk of CRS is assumed to be related to the mitigation of the peak of cytokine, considered as a PD endpoint [147, 38]. We propose a Bayesian approach incorporating PK/PD models in order to model the complex assumption between the probability of a PD-related toxicity and the dose regimen [64]. This method links the dose regimen and the PD profile using nonlinear mixed effects models and then the PD endpoint and the probability of toxicity using either a logistic or a hierarchical model. We observed good performance in selecting the MTD-regimen in a simulation study and the method is able to predict the probability of toxicity of dose regimens that were not tested in the trial.

The previous work focused on modeling the PD-related toxicity, however other potential DLT (DLT_o) should be included in the method to recommend the MTD-regimen. We propose to extend the previous method to include the modeling of DLT_o without PK/PD assumptions on the effect of the dose regimen [?]. We define the DLT as a bivariate binary endpoint to model the DLT_o while preserving the modeling of the first toxicity with PK/PD. We model the relationship between the

cumulative probability of DLT_o and the cumulative dose administered and develop three approaches to model the joint probability of DLT from previous models to account for potential association between toxicities: an independent model, a copula model and a conditional model. Our approaches can improve the performance of selection of the MTD-regimen in most simulated scenarios compared to the dose-escalation method implemented and can predict the DLT probability of new dose regimens that were not tested in the trial.

To conclude, our modeling proposal, to be applied at the end of the dose-escalation stage of the trial, can recommend the MTD-regimen and propose alternative dose regimens for expansion cohorts using additional information from PK/PD and various toxicity types.

Keywords: Dose regimens, pharmacokinetics and pharmacodynamics, Bayesian inference, early phase trials in oncology, bivariate toxicity.

Résumé

Titre : Modélisation Bayésienne de la toxicité pour l'évaluation de régimes de doses en essais de phase précoce en oncologie intégrant l'analyse pharmacocinétique et pharmacodynamique

Résumé : Les essais de phase I en oncologie cherchent à évaluer la toxicité du nouveau médicament et permettent d'affiner l'information sur le mécanisme d'action de la molécule, notamment par des études pharmacocinétiques et pharmacodynamiques (PK/PD). Leur objectif est de déterminer la dose maximale tolérée (MTD) définie comme la dose la plus élevée qui ne dépasse pas une certaine probabilité de toxicité dose-limitante (DLT) [42, 118]. La DLT est une toxicité binaire définie par l'occurrence d'effets indésirables graves. Des designs séquentiels d'escalade de dose ont été développés pour déterminer la MTD après un cycle de traitement, comme le 3+3 [144] ou la méthode de réévaluation continue (CRM) [119]. L'évaluation de régimes de doses, définis comme la combinaison du plan d'administration (fréquence et mode) et de la dose, peut permettre de réduire le risque de toxicité tout en maintenant une efficacité future. Certains designs ont été développés pour évaluer le plan d'administration en plus de la dose, mais ils ne permettent pas de considérer une hypothèse complexe sur la toxicité. Dans un essai de phase I (NCT03594955 [14]), des régimes de doses complexes, définis par une escalade de dose intra-patient, ont été mis en place afin de réduire le risque de la toxicité majeure attendue : le syndrome de libération des cytokines (cytokine release syndrome, CRS). Un design classique d'escalade de dose a été mis en place, mais il ne prend pas en compte l'information sur les régimes de doses. L'objectif de cette thèse est de développer une méthode, à appliquer à la fin de l'essai, afin de déterminer le régime de dose maximal toléré (MTD-régime).

La première partie de la thèse se place dans le contexte de l'essai lors duquel la réduction du risque de CRS est supposée être liée à la réduction du pic de cytokine, qui peut être considéré comme un critère PD [147, 38]. Nous proposons une méthode Bayésienne intégrant de la modélisation PK/PD afin de modéliser l'hypothèse complexe entre la probabilité d'une toxicité reliée à la PD (CRS) et le régime de doses [64]. Cette méthode modélise la relation entre le régime de doses et le profil PD à l'aide de modèles non linéaires à effets mixtes puis la relation entre le critère PD et la probabilité de toxicité à l'aide d'un modèle logistique ou hiérarchique. De bonnes performances de sélection du MTD-régime ont été observées sur une étude de simulations et la méthode est capable de prédire la probabilité de toxicité de régimes de doses qui n'ont pas été testés au cours de l'essai.

Lors du travail précédent, seule la toxicité majeure attendue est modélisée, or d'autres DLT potentielles (DLT_o) devraient également être prises en compte dans la recommandation du MTD-régime. Nous proposons d'étendre la méthode précédente

pour inclure la modélisation des DLT_o , sans hypothèse PK/PD sur le lien avec le régime de doses [?]. Nous définissons la DLT comme une variable bivariée binaire afin de modéliser la DLT_o en conservant la modélisation du premier type de toxicité avec de la PK/PD. Nous modélisons la probabilité cumulée de DLT_o avec la dose cumulée administrée et développons trois approches pour modéliser la probabilité jointe de DLT à partir des modèles précédents pour prendre en compte une éventuelle association entre les toxicités : un modèle indépendant, un modèle de copule et un modèle conditionnel. Nos approches permettent d'améliorer les performances de sélection du MTD-régime dans la plupart des scénarios de simulation par rapport à la méthode d'escalade de dose implémentée et peuvent également prédire la probabilité de DLT de nouveaux régimes de doses qui n'ont pas été testés au cours de l'étude.

En conclusion, notre approche de modélisation, à appliquer à la fin de la phase d'escalade de dose de l'essai, permet de recommander le MTD-régime et de proposer des régimes de doses alternatifs pour des cohortes d'expansion en utilisant de l'information additionnelle provenant de la PK/PD et de différents types de toxicité.

Mots clefs : Régimes de doses, modélisation pharmacocinétique et pharmacodynamique, inférence Bayésienne, essais de phase précoce en oncologie, toxicité bivariée.

Résumé long

Titre : Modélisation Bayésienne de la toxicité pour l'évaluation de régimes de doses en essais de phase précoce en oncologie intégrant l'analyse pharmacocinétique et pharmacodynamique

Contexte et objectifs de la thèse Les essais de phase I constituent la première étape du développement du médicament chez l'humain et ont lieu après les essais pré-cliniques, qui sont généralement effectués chez l'animal. L'objectif de ces essais est d'évaluer la tolérance du médicament et de déterminer la dose pour les futures phases du développement. Ces essais sont majoritairement effectués sur des volontaires sains, sauf pour certaines pathologies comme l'oncologie où la population est généralement composée de patients ayant reçu plusieurs lignes de traitements qui ont échoué.

Lors des essais de phase I en oncologie, les patients sont généralement traités lors de plusieurs cycles de traitement. Au cours d'un cycle, la dose ou le plan d'administration, constitué du mode d'administration (par exemple voie orale ou intraveineuse) et de la fréquence (par exemple une ou deux fois toutes les trois semaines), peuvent varier en cas de trop forte toxicité par exemple. La toxicité est mesurée par l'occurrence d'effets indésirables qui sont classés en différents grades selon leur sévérité, du grade 1 qui définit des toxicités mineures jusqu'au grade 5 qui est le décès [8]. Dans ces essais, tous les effets indésirables qui ont lieu au cours d'une fenêtre temporelle définie sont résumés en définissant une variable de toxicité binaire, la toxicité dose-limitante (DLT), qui définit des effets indésirables graves. La DLT est explicitement définie dans le protocole de chaque essai, elle peut par exemple être définie dès l'occurrence d'un effet indésirable de grade 3.

L'objectif des essais de phase I en oncologie est de trouver la dose maximale tolérée (MTD) définie comme la dose la plus élevée ayant une probabilité de DLT acceptable qui est en général la dose ayant la probabilité de DLT la plus proche d'une cible (par exemple 0.30) au cours du premier cycle de traitement [118]. Cette recherche de la MTD se base sur l'hypothèse des agents cytotoxiques qui suppose que la probabilité de toxicité et la probabilité d'efficacité augmentent de façon monotone avec la dose, ce qui signifie que la MTD a la probabilité d'efficacité la plus élevée parmi les doses avec une probabilité de toxicité acceptable comme représenté sur la Figure 6.1 [84]. Ces essais ont également pour objectif d'obtenir plus d'information sur le mécanisme d'action de la molécule, notamment par des analyses pharmacocinétiques (PK), qui étudient la relation entre la dose administrée et l'évolution de la concentration du médicament dans le corps avec le temps, et des analyses pharmacodynamiques (PD), qui étudient la relation entre la concentration du médicament et l'effet biochimique ou physiologique.

Les essais de phase I en oncologie sont soumis à de nombreuses contraintes




Figure 6.1 – Exemple de courbes de toxicité et d'efficacité

éthiques qui sont notamment d'inclure un petit nombre de patients, d'éviter d'attribuer des doses trop toxiques mais également d'éviter d'attribuer des doses trop faibles qui sont supposées être inefficaces dans la mesure où ces essais ont également un but thérapeutique pour les patients inclus [42]. Dans ce cadre, la randomisation des patients est donc inacceptable et des designs séquentiels d'escalade de dose ont été développés où les patients sont inclus par cohortes de petite taille et la dose attribuée à chaque cohorte dépend de la réponse observées sur les cohortes précédentes. Ces designs sont généralement classés en trois types, ceux basés sur un algorithme (comme le 3+3 [144]), ceux basés sur un modèle (comme la méthode par réévaluation continue, CRM [119]) et ceux assistés par un modèle qui combine la simplicité des designs algorithmiques avec la performance des designs basés sur un modèle [182].

Ces designs classiques ont été développés dans le contexte où le médicament est administré lors d'une seule administration, or, varier le régime de doses, défini par la combinaison du plan d'administration et de la dose, peut permettre d'améliorer la tolérance du produit tout en préservant une future efficacité. Certains designs ont été développés afin d'évaluer le plan d'administration en plus de la dose, mais ils ne permettent pas de prendre en compte une relation complexe entre la toxicité et le régime de doses qui peut être expliquée par le mécanisme d'action de la molécule [150, 160, 179]. Bien que les analyses PK/PD sont généralement des objectifs secondaires lors des essais de phase I en oncologie et peuvent évaluer l'effet de différents régimes de doses, peu de designs ont été développés pour incorporer ces données dans la détermination de la MTD [46, 155, 67].

Dans un essai de phase I en cours (NCT03594955 [14]) qui étudie un nouvel anticorps bispécifique engageant les lymphocytes T [15], la toxicité majeure attendue

est le syndrome de libération des cytokines (CRS) [137]. La CRS est supposée être liée au pic de cytokine [147], qui peut être considéré comme un marqueur PD, et il a été montré qu'administrer de façon répétée le produit peut permettre de diminuer le pic de cytokine [38]. D'autres toxicités (DLT_o) peuvent également avoir lieu, sans hypothèse a priori sur la relation avec le régime de doses via un marqueur PK/PD, et peuvent avoir une association avec la CRS [163, 132, 138, 60, 165]. Une escalade de dose intra-patient, fixée au début de l'essai, a donc été mise en place afin de diminuer le risque de CRS, où les patients reçoivent un régime de doses défini par une séquence de doses croissantes suivie par des administrations répétées de la dose finale reçue. Puisque les différents régimes de doses sont spécifiés avant le début de l'essai et sont supposés être ordonnés en terme de probabilité de toxicité, un design 3+3 a été choisi pour guider l'escalade de doses [27]. Or le 3+3, comme les designs classiques d'escalade de dose, ne permet pas de prendre en compte l'information des administrations multiples de différentes doses qui constituent le régime de doses. De plus, les designs développés pour optimiser le plan d'administration en plus de la dose ne prennent pas en compte des hypothèses complexes sur la probabilité de toxicité comme c'est le cas pour la CRS avec les administrations multiples.

L'objectif de ce travail de thèse est donc de développer une méthode afin de recommander le régime de dose maximal toléré (MTD-régime) qui permet de prendre en compte une hypothèse complexe sur la toxicité via un marqueur PD (la réduction du risque de CRS due à la mitigation du pic de cytokine dans le contexte de l'essai). Ce travail s'est effectué en deux parties. La première partie de la thèse s'est concentrée sur la modélisation d'une toxicité liée à un marqueur PD (la CRS dans le contexte de l'essai) qui a une relation complexe avec le régime de doses, qui incorpore de la modélisation PK/PD. La deuxième partie du travail s'est intéressée à l'inclusion des DLT_o dans la recommandation du MTD-régime, en conservant la modélisation incluant de la PK/PD, et en supposant une éventuelle association entre les deux types de toxicité. L'ensemble de la modélisation est représenté sur la Figure 6.2.

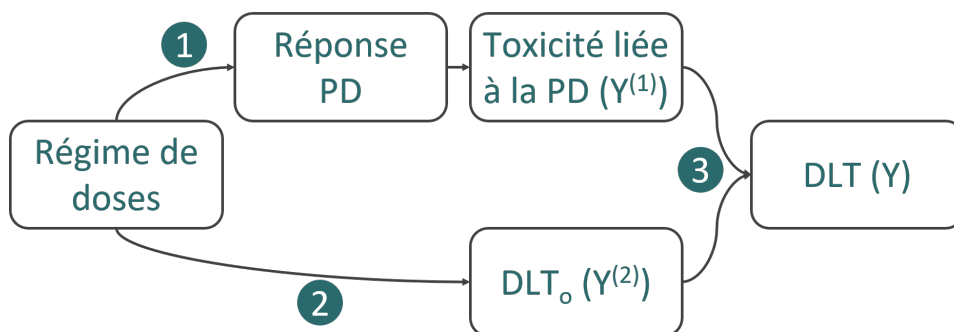


Figure 6.2 – Proposition de modélisation

Modélisation Bayésienne d'une toxicité avec le régime de doses qui incorpore de la modélisation PK/PD en essais de phase précoce en oncologie [64] Dans cette première partie, l'objectif est de modéliser une relation complexe entre une toxicité et le régime de doses en incorporant de la modélisation PK/PD. Ce travail s'inspire de l'essai donné en exemple lors duquel l'escalade de dose intra-patient permet de diminuer le risque de CRS par la mitigation du pic de

cytokine, considéré comme un marqueur PD. Nous proposons de modéliser la relation entre le régime de doses et la probabilité de toxicité en modélisant la relation entre le régime de doses et le profil PD, puis la relation entre le marqueur PD et la probabilité de toxicité.

Nous modélisons la relation entre le régime de doses et le profil PD à l'aide de modèles non linéaires à effets mixtes. Le modèle PK modélise la relation entre le régime de doses et la concentration de médicament et le modèle PD modélise la relation entre la concentration et le profil PD (la réponse en cytokine dans le contexte de l'essai). Ces modèles permettent d'incorporer de la connaissance biologique et d'inclure de la variabilité entre les patients, qui est supposée être élevée dans ce type de données. Dans le contexte de l'essai, ces modèles permettent de modéliser la mitigation du pic de cytokine avec l'escalade de dose intra-patient comme représenté sur la Figure 6.3.

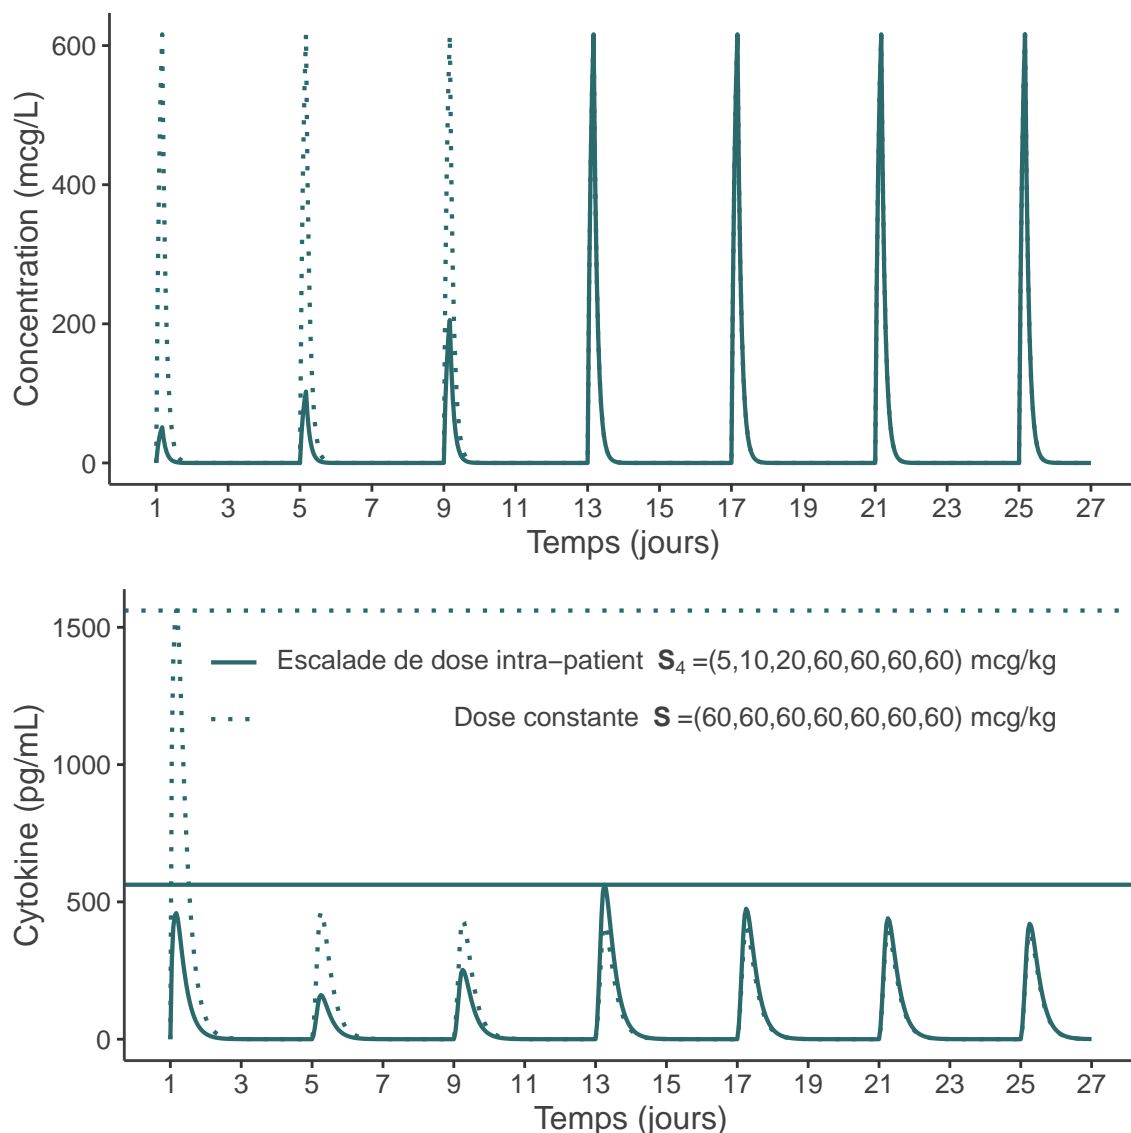


Figure 6.3 – Effet de l'escalade de dose intra-patient sur la réponse PK/PD

Pour modéliser la relation entre le marqueur PD et la probabilité de toxicité, nous proposons deux approches Bayésiennes : un modèle logistique et un modèle

hiérarchique. Dans ces deux modèles, nous incluons une valeur de population pour le marqueur PD ce qui nous permet de définir les distributions a priori à partir des probabilités de toxicité de chaque régime de doses que l'on imagine au début de l'essai. Nous mesurons l'information apportée par la distribution a priori avec l'*effective sample size* (ESS) [109, 175] qui représente le nombre fictif de patients apportés par cette distribution.

Enfin, pour estimer la probabilité de toxicité de chaque régime de doses en incluant la variabilité inter-individuelle sur la PK/PD, nous proposons une méthode de Monte-Carlo. Nous définissons alors le MTD-régime comme le régime de doses ayant la probabilité de toxicité la plus proche de la cible.

Comme notre méthode nécessite que les données PK/PD soient collectées et que les modèles PK/PD soient développés et fiables, nous avons développé notre approche pour qu'elle soit appliquée à la fin de la phase d'escalade de dose de l'essai, ce qui nécessite donc de définir un design afin de guider l'escalade de dose.

Nous avons évalué les performances de notre méthode sur une étude de simulation, où les modèles PK/PD sont inspirés du blinatumomab, qui est un autre anticorps bispécifique engageant les lymphocytes T [2, 184, 38]. Nous avons défini plusieurs scénarios de simulation afin de varier la position du vrai MTD-régime, le design d'escalade de dose implémenté (3+3 ou CRM modifiée) et le type de régimes de doses étudiés. Nous avons observé de bonnes performances de sélection du vrai MTD-régime par rapport au design d'escalade de dose implémenté. Nous avons également observé de moins bons résultats après un 3+3 qu'une CRM modifiée ce qui indique que le design d'escalade de dose implémenté a un impact sur les performances de notre méthode et nous recommandons un design plus complexe que le 3+3 afin d'avoir une meilleure qualité de données. De plus, vu que notre méthode modélise la relation entre le régime de doses administré et la probabilité de toxicité, elle peut prédire la probabilité de toxicité de nouveaux régimes de doses qui n'ont pas été évalués au cours de l'essai. Notre méthode peut donc permettre de proposer différents régimes de doses à la fin de l'essai, qui peuvent être évalués sur des cohortes d'expansion afin de comparer leur efficacité par exemple.

En conclusion, nous avons proposé une approche performante pour recommander le MTD-régime qui incorpore de la modélisation PK/PD afin de prendre en compte une hypothèse complexe sur la relation entre une toxicité et le régime de doses. La première limite de notre approche est qu'elle doit être appliquée à la fin de l'étape d'escalade de dose, pour que les données PK/PD soient collectées et les modèles PK/PD soient développés, et donc le design d'escalade de dose implémenté peut impacter les résultats. Notre méthode nécessite également de fortes hypothèses sur le mécanisme générant la toxicité, mais elle peut être étendue à d'autres critères PK/PD.

Modélisation Bayésienne d'une toxicité bivariée avec le régime de doses en essais de phase précoce en oncologie [?] Dans le travail précédent, nous avons proposé une méthode pour recommander le MTD-régime qui modélise un certain type de toxicité ayant une relation complexe avec le régime de doses en incorporant de la modélisation PK/PD. Or d'autres toxicités, appelées DLT_o , peuvent également avoir lieu et doivent être prises en compte dans la recommandation du MTD-régime. Étant donné que le mécanisme d'action générant les DLT_o n'est pas connu contrairement au premier type de toxicité et n'est pas a priori lié à un

critère PK/PD, nous proposons de modéliser la DLT, Y , comme une variable bivariable binaire définie par $Y = \max(Y^{(1)}, Y^{(2)})$ où $Y^{(1)}$ est le premier type de toxicité supposée être reliée au critère PD, et $Y^{(2)}$ est la DLT_o. $Y^{(1)}$ est modélisée avec de la PK/PD comme présenté dans le premier travail. Pour $Y^{(2)}$, comme le mécanisme générant la toxicité n'est pas connu, nous proposons de modéliser la probabilité cumulée de toxicité avec la dose cumulée administrée afin de prendre en compte les administrations multiples [154]. Comme pour le premier travail, nous définissons les distributions a priori à partir des probabilités de DLT de chaque régime de doses que l'on imagine au début de l'essai, en supposant que $Y^{(1)}$ et $Y^{(2)}$ sont indépendantes avec des probabilités de toxicité égales. Nous mesurons l'information apportée par la distribution a priori avec l'ESS. Enfin, nous proposons trois méthodes jointes pour modéliser la probabilité de DLT, en supposant que les deux types de toxicité, $Y^{(1)}$ et $Y^{(2)}$, peuvent être associés. Nous proposons d'abord un modèle indépendant en supposant que les deux types de toxicité sont indépendants. Nous proposons ensuite un modèle de copule qui permet de modéliser la probabilité de DLT à partir des probabilités marginales de chaque type de toxicité et d'un paramètre d'association. Finalement, puisque nous supposons que les deux types de toxicité n'ont pas lieu au même moment au cours de l'administration, c'est-à-dire que la toxicité liée au critère PD a tendance à avoir lieu au début des administrations alors que les DLT_o ont lieu à la fin du régime de doses, nous proposons de modéliser la probabilité de Y à partir de la probabilité marginale de $Y^{(1)}$ et de la probabilité conditionnelle de $Y^{(2)}$ sachant que l'on n'a pas observé $Y^{(1)}$. Cette méthode permet de prendre en compte l'association entre les deux types de toxicité, sans la modéliser explicitement.

Nous avons évalué les performances de nos trois approches jointes sur une étude de simulation où les modèles PK/PD sont inspirés du blinatumomab. Pour simuler les DLT_o, nous définissons les probabilités de DLT_o à chaque administration conditionnellement au fait qu'il n'y a pas eu de DLT_o lors des administrations précédentes et conditionnellement au fait que le premier type de toxicité a lieu ou non afin de définir l'association entre les deux types de toxicité. Nous mesurons cette association entre les deux types de toxicité par le risque relatif qui représente la probabilité d'observer une DLT_o sachant que le premier type de toxicité aura lieu. Nous avons défini plusieurs scénarios de simulation afin d'évaluer l'effet de la position du vrai MTD-régime, l'association entre les deux types de toxicité et la proportion de chaque type de toxicité et nous appliquons nos méthodes jointes à la fin de l'essai simulé sous une CRM modifiée. Nos méthodes jointes ont de meilleures performances de sélection du vrai MTD-régime par rapport à la CRM dans la majorité des scénarios. En cas de faible association entre les deux types de toxicité, les trois méthodes jointes ont des résultats similaires. Or, quand l'association augmente, le modèle de copule et le modèle conditionnel semblent avoir de meilleurs résultats que le modèle indépendant puisque qu'ils prennent en compte l'association entre les deux types de toxicité. Nous avons également observé que le modèle indépendant et le modèle de copule ont tendance à sous estimer la probabilité marginale de DLT_o. En effet, vu que l'on suppose que les toxicités liées à la réponse PD ont plutôt tendance à avoir lieu au début du régime de doses et que l'administration du médicament est arrêtée si une DLT a lieu, on n'observe les DLT_o que lorsque les toxicités reliées à la PD n'ont pas lieu, sauf dans le cas où les deux toxicités arrivent après la même administration. Enfin, nos trois méthodes jointes modélisent la relation entre le régime de doses et la probabilité de DLT et peuvent donc prédire la probabilité de DLT

de régimes de doses qui n'ont pas été testés au cours de l'essai, mais qui pourraient avoir une pertinence clinique, au niveau de l'efficacité par exemple.

En conclusion, nous avons proposé une approche qui permet de modéliser explicitement une toxicité d'intérêt ayant une relation complexe avec le régime de doses par de la modélisation PK/PD tout en prenant en comptes d'autres toxicités qui peuvent avoir lieu et être associées au premier type de toxicité. Au niveau des limites, notre approche a été développée pour être appliquée à la fin de la phase d'escalade de dose de l'essai par rapport à la disponibilité des données PK/PD mais également car elle nécessite d'avoir assez de données collectées pour effectuer les différentes étapes de modélisation. Nous avons également observé que le paramètre du modèle de copule était difficile à estimer, même en cas de forte association entre les deux types de toxicité. Enfin, nous avons considéré un modèle simple sur les DLT_o car il est possible qu'un faible nombre de ce type de toxicité soit observé et peu d'hypothèses sont disponibles contrairement au premier type de toxicité.

Perspectives La modélisation du premier type de toxicité s'est concentrée sur un critère PD (le pic de cytokine dans le contexte de l'essai donné en exemple), mais notre approche de modélisation peut être étendue à d'autres critères. Une extension de ce travail pourrait être de modéliser la DLT_o avec un critère PK (par exemple l'AUC) ce qui permettrait d'intégrer directement la sensibilité du patient au médicament. Dans notre travail, la recommandation du MTD-régime est basée sur l'estimation de la probabilité de toxicité, or une extension de ce travail serait de baser cette recommandation sur des intervalles de toxicité afin de prendre en compte toute l'information sur la distribution a posteriori. Enfin, une dernière extension de ce travail pourrait être de considérer des données de temps à l'événement à la place de données longitudinales.

Appendix A

Bayesian inference

A.1 Bayesian paradigm

Bayesian and frequentist approaches have different philosophies. In the Bayesian approach, all unknowns have probability distributions while in the frequentist approach, the probability is only defined for the data. In particular, for Bayesians, parameters are random unobserved variables while for frequentists, parameters are fixed unknowns quantities [24].

Let y be the observed data and θ the parameter. Bayesian inference is based on the property of conditional distributions, known as Bayes' rule [63]:

$$p(\theta, y) = p(\theta|y)p(y) = p(y|\theta)p(\theta) \quad (\text{A.1})$$

Equation A.1 leads to

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)} \quad (\text{A.2})$$

- $p(y|\theta)$ can be seen as a function of θ with y fixed which is then is the likelihood of observations defined with a probability model
- $p(\theta)$ is the prior distribution that summarizes all information on the parameter before observing the data
- $p(y) = \int_{\Theta} p(y|\theta)p(\theta) d\theta$ is the marginal distribution of data
- $p(\theta|y)$ is the posterior distribution that updates the information on the parameters after collecting the data

For example, let's assume that we aim to estimate the probability of response of a binary endpoint in a cohort, and observe $y = 5$ patients who respond in a cohort of 10 patients. We can model the number of responses in the cohort using a binomial model and calculate the likelihood represented in Figure A.1. For the prior, we can choose a beta distribution and define the parameters based on the amount of initial knowledge. Figure A.1 illustrates three choices of prior information, from the least informative one in the left part to the most informative one in the right part. In case of a non informative prior (left part), the posterior distribution is very close

to the likelihood. In case of a quite informative prior (middle part), the posterior distribution is influenced by both the prior and the likelihood. In case of a very informative prior, the posterior is mostly influenced by the prior distribution.

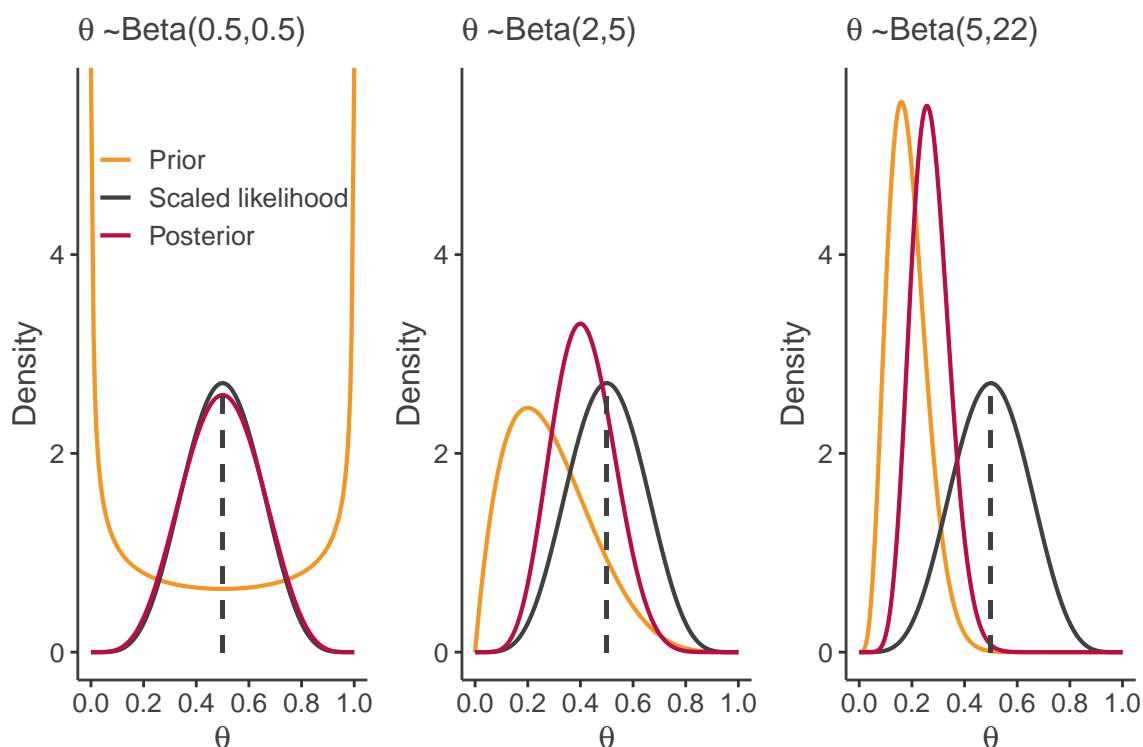


Figure A.1 – Illustration of Bayes’ rule with the beta-binomial model for various prior distributions. The dashed vertical line represents the maximum likelihood estimator.

The choice of the prior distribution is therefore crucial in Bayesian inference and should be carefully selected as it can have a strong impact on the inference. In the context of clinical trials, the prior distribution can be defined from clinicians’ opinion or from previous data, for example from previous trials of the same molecule in the same or a different population or from trials of a similar molecule.

Bayesian inference is mainly based on the posterior distribution from which several point estimates can be defined as the mean or maximum a posteriori and the credible interval can be defined but the entire posterior distribution can also be used for the inference.

A.2 Computation of the posterior distribution

The main challenge for Bayesian methods is the computation of the posterior distribution. In simple cases, the posterior distribution can be computed in a closed form. A convenient case is conjugacy meaning that the posterior distribution follows the same parametric form as the prior distribution (for example the beta-binomial model). However, for more complex models, Markov chain Monte Carlo (MCMC) algorithms have been developed to sample from the posterior distributions. The idea of MCMC is to simulate a Markov Chain whose stationary distribution is the posterior distribution. In practice, at each iteration t of the algorithm, parameter $\theta^{(t)}$ is sampled from the previous sampled value $\theta^{(t-1)}$, and for a large number of iterations, the distribution of the current sampled values should be close to the stationary distribution. Multiple MCMC algorithms have been developed [63].

Gibbs sampler: This algorithm assumes that samples can be generated from each of their conditional distributions. Let $\theta = (\theta_1, \dots, \theta_K)$.

For each iteration t :

- Draw $\theta_1^{(t)}$ from $p\left(\theta_1|\theta_2^{(t-1)}, \theta_3^{(t-1)}, \dots, \theta_K^{(t-1)}, y\right)$
- Draw $\theta_2^{(t)}$ from $p\left(\theta_2|\theta_1^{(t)}, \theta_3^{(t-1)}, \dots, \theta_K^{(t-1)}, y\right)$
- ...
- Draw $\theta_K^{(t)}$ from $p\left(\theta_K|\theta_1^{(t)}, \theta_2^{(t)}, \dots, \theta_{K-1}^{(t)}, y\right)$

This algorithm can be used for conditionally conjugates models.

Metropolis-Hastings algorithm: This rejection algorithm can be used when we cannot sample directly from conditional distributions and is based on a proposal distribution g .

For each iteration t :

- Draw θ^* from a proposal distribution $g(\theta^*|\theta^{(t-1)})$
- Calculate $r = \frac{p(\theta^*|y)g(\theta^{(t-1)}|\theta^*)}{p(\theta^{(t-1)}|y)g(\theta^*|\theta^{(t-1)})}$
- Set $\theta^t = \begin{cases} \theta^* & \text{with probability } \min(1, r) \\ \theta^{t-1} & \text{otherwise} \end{cases}$

The Metropolis-Hastings algorithm can be implemented when p is specified up to a constant. This feature is particularly useful in Bayesian inference where the posterior distribution is proportional to the product of the likelihood and the prior, $p(\theta|y) \propto p(y|\theta)p(\theta)$. The Gibbs sampler and the Metropolis-Hastings algorithm can be combined in case we can directly sample from some conditional distributions, and not from others.

Hamiltonian Monte Carlo: The idea of Hamiltonian Monte Carlo comes from physics and the algorithm aims to explore more efficiently the target distribution using information on its geometry by defining the gradient of the target distribution and auxiliary momentum parameters [26]. This algorithm is implemented in stan software [16]. Betancourt [26] explains the concept of Hamiltonian Monte Carlo and proposes an analogy in physics where a planet represents the mode of the target distribution, the planet’s gravitational field represents the gradient of the target density and the space around the planet where we want an object (as a satellite) to orbit represents the set to explore. To avoid the satellite to crash in the planet or the satellite to be lost in space, we have to add enough, but not too much, momentum for the satellite to orbit.

Let $p(\theta)$ be the distribution to sample (is our case, this corresponds to the posterior distribution $p(\theta|y)$). Let ρ be the auxiliary momentum variables that usually follow a multivariate normal distribution $\rho \sim \mathcal{N}(0, M)$.

The joint distribution $p(\rho, \theta) = p(\rho|\theta)p(\theta)$ defines a Hamiltonian as follows:

$$H(\rho, \theta) = T(\rho|\theta) + V(\theta) \tag{A.3}$$

where $T(\rho|\theta) = -\log(p(\rho|\theta))$ is the *kinetic energy* and $V(\theta) = -\log(p(\theta))$ is the *potential energy*.

For each iteration t :

- Draw ρ from its distribution $\mathcal{N}(0, M)$ and set $\theta = \theta^{(t-1)}$
- Solve the Hamiltonian system $\begin{cases} \frac{d\theta}{dt} = +\frac{\partial H}{\partial \rho} = +\frac{\partial T}{\partial \rho} \\ \frac{d\rho}{dt} = -\frac{\partial H}{\partial \theta} = -\frac{\partial V}{\partial \theta} \end{cases}$, assuming that the momentum density is independent from the target density. This system can be numerically solved by L leapfrog steps where ϵ is a small time interval:

- $\rho \leftarrow \rho - \frac{\epsilon}{2} \frac{\partial V}{\partial \theta}$
- $\theta \leftarrow \theta + \epsilon M^{-1} \rho$
- $\rho \leftarrow \rho - \frac{\epsilon}{2} \frac{\partial V}{\partial \theta}$

- Metropolis acceptance step for the state at the end of the leapfrog algorithm (ρ^*, θ^*) . Calculate $r = \exp(H(\rho^{(t-1)}, \theta^{(t-1)}) - H(\rho^*, \theta^*))$.

$$\text{Set } \theta^{(t)} = \begin{cases} \theta^* & \text{with probability } \min(1, r) \\ \theta^{(t-1)} & \text{otherwise} \end{cases}$$

Parameters M , ϵ and L can be tuned in stan, in particular the number of leapfrog steps can be dynamically adapted with the no-U-turn sampling algorithm.

Appendix B

PK/PD models and estimation

B.1 Details on the PK/PD models

Blinatumomab is a BiTE that binds to both CD3 expressed on T cells and CD19 expressed on B cells [2]. It was initially approved by the FDA for the treatment of Philadelphia chromosome-negative relapsed or refractory acute lymphoblastic leukemia, was granted accelerated approval in B-cell precursor acute lymphoblastic leukemia [3].

Zhu et al. [184] performed a population PK analysis with nonlinear mixed effects models using the data from 4 studies of blinatumomab administered in continuous IV in various indications with doses ranging from 0.5 to 90 $\mu\text{g}/\text{m}^2/\text{day}$ or 9 to 28 $\mu\text{g}/\text{day}$. They selected a one compartment linear model where parameters were lognormally distributed. Estimated parameters are defined in Table B.1, where the clearance parameter was defined from one subpopulation as they found a bimodal distribution for the clearance parameter. They also found that the cytokines values increased quickly with the dose after the infusion of blinatumomab during the first week of therapy, among which IL-10, IL-6, and IFN- γ .

Chen et al. [38] developed a PD model for the cytokine response of blinatumomab from a phase I trial in patients with relapsed non-Hodgkin lymphoma. Cytokine profiles were available after either a step dose of 5 to 15 $\mu\text{g}/\text{m}^2/\text{day}$ or a single dose of 60 $\mu\text{g}/\text{m}^2/\text{day}$. To model cytokine mitigation with multiple administrations, they assumed that cytokine production was stimulated by the release, RL, but inhibited by the negative feedback, IH, as illustrated in Figure B.1.




Figure B.1 – Cytokine model.

The cytokine model is defined as follows:

$$\frac{dR(t)}{dt} = \text{RL} (1 - \text{IH}) - k_{\text{deg}}R(t), \quad (\text{B.1})$$

where R is the cytokine response, RL is the release, IH is the negative feedback and k_{deg} is the first order degradation rate.

The cytokine release is stimulated by the synapse modeled as follows:

$$RL(t) = \frac{E_{\text{max}}\text{Syn}(t)^H}{EC_{50}^H + \text{Syn}(t)^H}, \quad (\text{B.2})$$

where Syn is the synapse concentration, E_{max} is the maximum cytokine release rate, EC_{50} is the synapse concentration to achieve 50% of maximum cytokine release rate and H is the Hill coefficient.

For blinatumomab, synapse concentration is impacted by tumor cell kinetics as follows:

$$\text{Syn}(t) = C(t) * \text{Tu}_{\text{tot}}(t), \quad (\text{B.3})$$

where C is drug concentration and

$$\frac{d\text{Tu}_{\text{tot}}(t)}{dt} = -k_{\text{kill}}\text{Tu}_{\text{tot}}(t) \quad (\text{B.4})$$

To simplify the PD model in our work, we assumed that synapse concentration is mainly based on drug concentration, as for solid tumors, meaning that

$$\text{Syn}(t) = C(t) \quad (\text{B.5})$$

Finally, cytokine production is inhibited by the negative feedback parameter defined from the AUC of the cytokine response, AUC_R as follows:

$$IH(t) = \frac{I_{\text{max}}\text{AUC}_R(t)}{\frac{IC_{50}}{K^{N-1}} + \text{AUC}_R(t)}, \quad (\text{B.6})$$

where I_{max} is the maximum inhibition, IC_{50} is the cumulative cytokine exposure to achieve 50% of the maximum inhibition, K is the priming factor and N is the number of doses. Chen et al. [38] focused on the first two doses, meaning that $N=2$.

For our simulations, we therefore considered the following PD model for the cytokine response:

$$\frac{dR(t)}{dt} = \frac{E_{\text{max}}C(t)^H}{EC_{50}^H + C(t)^H} \left(1 - \frac{I_{\text{max}}\text{AUC}_R(t)}{\frac{IC_{50}}{K^{N-1}} + \text{AUC}_R(t)} \right) - k_{\text{deg}}R(t) \quad (\text{B.7})$$

The PK parameters estimated by Zhu et al. [184] and the PD parameters estimated by Chen et al. [38] are displayed in Table B.1.

B.1. DETAILS ON THE PK/PD MODELS

	Parameter	Estimated value [184, 38] (%CV)	Value for simulations (%CV)	Unit	Description
PK model	Cl	1.36 (41.9)	1.36 (41.9)	L/h	Clearance of elimination
	V	3.40 (0)	3.40 (0)	L	Volume of distribution
PD model	E_{\max}	$3.59 \cdot 10^3$ (14)	$3.59 \cdot 10^5$ (14)	pg/mL/h	Maximum cytokine release rate
	EC_{50}	$1.00 \cdot 10^4$ (0)	$1.00 \cdot 10^4$ (0)	ng/mL	Drug exposure for half-maximum cytokine release rate
	H	$9.20 \cdot 10^{-1}$ (3)	$9.20 \cdot 10^{-1}$ (3)		Hill coefficient for cytokine release
	I_{\max}	1 (0)	0.995 (0)		Maximum inhibition of cytokine release
	IC_{50}	1.82×10^2 (12)	1.82×10^4 (12)	pg/mL/h	Cytokine exposure for half-maximum cytokine inhibition
	k_{deg}	$1.80 \cdot 10^{-1}$ (13)	$1.80 \cdot 10^{-1}$ (13)	h^{-1}	Degradation rate for cytokine
	K	2.83 (36)	2.83 (36)		Priming factor for cytokine release

Table B.1 – Estimated PK/PD parameters for blinatumomab [184, 38] and PK/PD parameters used for the simulations.

We considered some modifications of the PD parameters for our simulation studies, as shown in Table B.1, to mimic our assumptions on the motivating trial. To define the changes, we can show the PK/PD profiles obtained for the dose regimens given in Section 3.1 and displayed in Table B.2.

	Day 1	Day 5	Day 9	Day 13	Day 17	Day 21	Day 25
S_1	1	5	10	10	10	10	10
S_2	1	5	10	20	20	20	20
S_3	5	10	20	40	40	40	40
S_4	5	10	20	60	60	60	60
S_5	10	20	60	100	100	100	100
S_6	10	20	60	120	120	120	120

Table B.2 – Example of a panel of dose regimens defined with intra-patient dose-escalation (doses in $\mu\text{g}/\text{kg}$).

Population PK profiles obtained after the first 5 administrations of these dose regimens (assuming 4 hours of infusion) with the parameters estimated by Zhu et al. [184] are shown in Figure B.2.




Figure B.2 – Concentration profiles from Zhu et al. [184].

The cytokine profiles using the model defined in Equation B.7 after the first 5 administrations of these dose regimens with the parameters estimated by Chen et al. [38], are shown in Figure B.3. With these parameters, the peak of cytokine occurs during the first administration. To allow more flexibility in the timing of occurrence of the peak of cytokine, we modified the value of the I_{\max} parameter and we increased the values of E_{\max} and IC_{50} to increase the scale of the cytokine response. With these new parameters, the population cytokine profiles are illustrated in Figure B.4.




Figure B.3 – Cytokine profiles from the model defined in Equation B.7 using the initial parameters from Chen et al. [38].




Figure B.4 – Cytokine profiles from the model defined in Equation B.7 using the modified parameters.

B.2 Estimation of population parameters for non-linear mixed effects models

In this part, a general overview of estimation methods for nonlinear mixed effects models is provided, based on [17, 45].

In the context of nonlinear mixed effects models, we consider the model introduced in Section 2.1.3 defined as follows:

$$y_{i,j} = f(\boldsymbol{\theta}_i, \mathbf{x}_{ij}) + g(\boldsymbol{\theta}_i, \mathbf{x}_{i,j}, \boldsymbol{\xi}) \epsilon_{ij}, \quad \epsilon_{ij} \sim \mathcal{N}(0, 1), \quad (\text{B.8})$$

where f and g are nonlinear functions, and $\boldsymbol{\theta}_i$ are the patient's individual parameters defined from the individual random effects $\boldsymbol{\eta}_i$. Let $\boldsymbol{\Psi}$ denote the population parameters that can be estimated by maximum likelihood, where the likelihood is defined as follows:

$$\mathcal{L}(\boldsymbol{\Psi}; \mathbf{y}) = \prod_{i=1}^n \int p(\mathbf{y}_i | \boldsymbol{\eta}_i, \boldsymbol{\Psi}) p(\boldsymbol{\eta}_i, \boldsymbol{\Psi}) d\boldsymbol{\eta}_i \quad (\text{B.9})$$

As this likelihood does not have an analytical expression, first approaches proposed to estimate the likelihood, for example using its first order approximation (FO or FOCE) [136, 98]. With these approximations, the maximum likelihood estimate can be obtained with a standard algorithm, for example the Newton-Raphson iterative algorithm. But the main drawback of these methods are the approximation of the likelihood. Furthermore, the Newton-Raphson can converge to a local minimum/maximum.

The Expectation-Maximization (EM) algorithm is an alternative iterative method to estimate the maximum likelihood, which was developed for partially observed data. In the case of nonlinear mixed effects models, the individual random effects $\boldsymbol{\eta}_i$ can be considered as missing data. Each iteration k of the EM algorithm updates the estimation from $\boldsymbol{\Psi}_{k-1}$ to $\boldsymbol{\Psi}_k$ and is divided in two steps:

- **Expectation (E):** Evaluation of the conditional expected value of the log likelihood since the individual random effects are not observed

$$Q(\boldsymbol{\Psi} | \mathbf{y}, \boldsymbol{\Psi}_k) = \mathbb{E}[\log(\mathcal{L}(\boldsymbol{\Psi}; \mathbf{y}, \boldsymbol{\eta})) | \mathbf{y}, \boldsymbol{\Psi}_k] \quad (\text{B.10})$$

- **Maximization (M):** $\boldsymbol{\Psi}_{k+1}$ defined by maximizing previous quantity

$$\boldsymbol{\Psi}_{k+1} = \arg \max_{\boldsymbol{\Psi}} (Q(\boldsymbol{\Psi} | \mathbf{y}, \boldsymbol{\Psi}_k)) \quad (\text{B.11})$$

However, with nonlinear mixed effects models, the quantity defined in step (E) cannot be computed in a closed form, but can be approximated. For example, in the Monte Carlo EM algorithm [164], the Q function in step (E) is approximated by Monte Carlo simulations. The Stochastic Approximation EM (SAEM) [52] replaces the (E) step by two steps at each iteration k :

- **Simulation (S) :** The random effects for all individuals ($\boldsymbol{\eta}^{(k)}$) are simulated from the conditional distribution $p(\boldsymbol{\eta} | \mathbf{y}, \boldsymbol{\Psi}_{k-1})$. An MCMC algorithm can be used when $\boldsymbol{\eta}^{(k)}$ cannot be simulated exactly under the conditional distribution [82].

- **Stochastic approximation (SA) :**

$$Q(\Psi|\mathbf{y}, \Psi_k) = Q(\Psi|\mathbf{y}, \Psi_{k-1}) + \gamma_k \left(\log \left(\mathcal{L}(\Psi; \mathbf{y}, \boldsymbol{\eta}^{(k)}) \right) - Q(\Psi|\mathbf{y}, \Psi_{k-1}) \right) \quad (\text{B.12})$$

where γ is a positive decreasing sequence

The SAEM algorithm therefore calculates the maximum likelihood estimate without computing the likelihood, which can be calculated at the end of the algorithm. SAEM was shown to be efficient and is implemented in monolix software [7].

Appendix C

Supporting information: Bayesian
dose regimen assessment in early
phase oncology incorporating
pharmacokinetics and
pharmacodynamics

Supporting information for Bayesian dose-regimen assessment in early phase oncology incorporating pharmacokinetics and pharmacodynamics by Emma Gerard, Sarah Zohar, Hoai-Thu Thai, Christelle Lorenzato, Marie-Karelle Riviere and Moreno Ursino

Web Appendix A: Toxicity probability for the hierarchical-DRtox

The prediction of the toxicity probability of a new patient having r_i^M as summary PD endpoint for the hierarchical-DRtox can be obtained as:

$$\begin{aligned}
\mathbb{P}(Y_i = 1) &= 1 - \mathbb{P}(Y_i = 0) \\
&= 1 - \mathbb{P}(Y_{i,1} = 0, \dots, Y_{i,j_i} = 0) \\
&= 1 - \mathbb{P}\left\{Z_i > \log\left(\frac{r_{i,1}}{\bar{r}_{k_{50}}^M}\right), \dots, Z_i > \log\left(\frac{r_{i,j}}{\bar{r}_{k_{50}}^M}\right)\right\} \\
&= 1 - \mathbb{P}\left\{Z_i > \log\left(\frac{r_i^M}{\bar{r}_{k_{50}}^M}\right)\right\} \\
&= F_z\left\{\log\left(\frac{r_i^M}{\bar{r}_{k_{50}}^M}\right)\right\}
\end{aligned} \tag{1}$$

Web Appendix B: Definition of the PK/PD models

For the PK model, we considered a one-compartment infusion model. Let $T_{inf}^{(j)}$ be the duration of the infusion of the j^{th} administration. We assumed that the delay between successive doses was greater than infusion duration, meaning that $t_{j+1} - t_j > T_{inf_j}$ for $j \in \{1, \dots, J-1\}$. Let V be the distribution volume, Cl the clearance of elimination and k the micro-constant defined as $k = \frac{Cl}{V}$. The concentration is then computed as:

$$\begin{aligned}
C(t) &= \sum_{j=1}^J \mathbb{1}_{\{t-t_j > T_{inf_j}\}} \frac{d_{k,j}}{T_{inf_j}} \frac{1}{kV} \left(1 - e^{-kT_{inf_j}}\right) e^{-k(t-t_j-T_{inf_j})} + \\
&\quad \mathbb{1}_{\{t-t_j \leq T_{inf_j} \ \& \ t-t_j \geq 0\}} \frac{d_{k,j}}{T_{inf_j}} \frac{1}{kV} \left\{1 - e^{-k(t-t_j)}\right\}
\end{aligned} \tag{2}$$

For the PD model developed by Chen et al. (2019), the cytokine profile is the result of cytokine release (RL), inhibition effect on cytokine release (IH) and first order degradation (k_{deg}):

$$\frac{dE(t)}{dt} = RL(1 - IH) - k_{deg}E(t) \tag{3}$$

They modeled that cytokine is released due to the synapse created by the binding of the drug to both T-cells and tumor cells as

$$RL = \frac{E_{max}Syn(t)^H}{EC_{50}^H + C(t)^H} \tag{4}$$

where $Syn(t)$ is the synapse concentration. To simplify, we assumed that $Syn(t) = C(t)$ meaning that only longer treatment will impact tumor burden as for solid tumors. E_{max} represents the maximum cytokine release rate and EC_{50} represents the synapse (or drug for our simplification) to achieve 50% of maximum cytokine release rate. The Hill coefficient H represents the slope of the exposure-response relationship.

Finally they modeled the inhibition effect as

$$IH = \frac{I_{max}AUC_E(t)}{\frac{IC_{50}}{K^{J-1}} + AUC_E(t)} \quad (5)$$

where $AUC_E(t) = \int_{x=0}^t E(x) dx$, I_{max} is the maximum inhibition, IC_{50} is the cumulative cytokine exposure to achieve 50% of the maximum inhibition, K is the priming factor upon repeating dosing and J is the number of doses of the regimens.

Web Appendix C: Definition of the dose-regimens

Table 1: Panel of dose-regimens from S_1 to S_6 for the three main scenarios (in mcg/kg).

		Day 1	Day 5	Day 9	Day 13	Day 17	Day 21	Day 25
Scenario 1	S_1	1	5	10	10	10	10	10
	S_2	1	5	10	20	20	20	20
	S_3	1	5	10	25	25	25	25
	S_4	5	10	25	50	50	50	50
	S_5	10	25	50	75	75	75	75
	S_6	10	25	50	100	100	100	100
Scenario 2	S_1	1	5	10	25	25	25	25
	S_2	5	10	25	50	50	50	50
	S_3	10	25	50	75	75	75	75
	S_4	10	25	50	100	100	100	100
	S_5	10	25	50	100	150	150	150
	S_6	25	50	100	150	200	200	200
Scenario 3	S_1	1	5	10	20	30	40	40
	S_2	1	5	10	20	40	40	40
	S_3	5	10	20	40	40	40	40
	S_4	10	20	40	40	40	40	40
	S_5	20	40	40	40	40	40	40
	S_6	40	40	40	40	40	40	40

Web Appendix D: Flowchart of the 3+3 escalation scheme

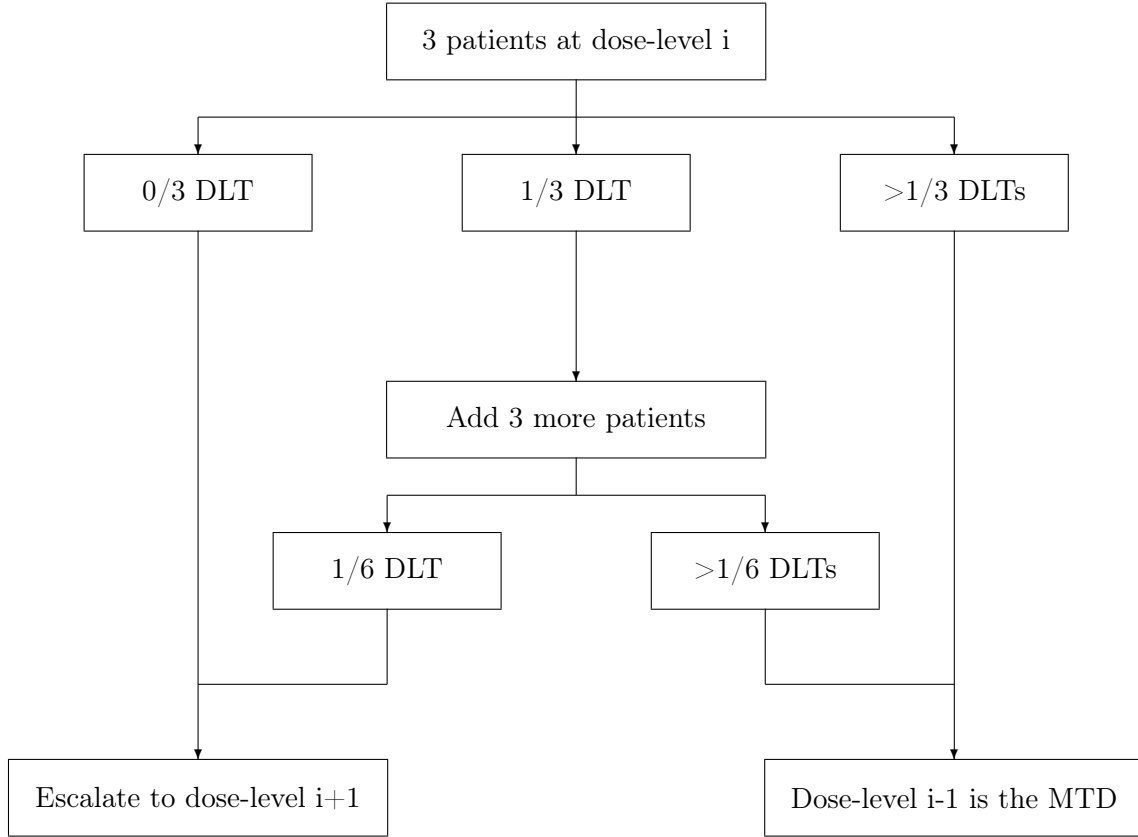


Figure 1: Flowchart of the 3+3 escalation scheme inspired from Chevret et al. (2006), the 3+3 design is a special case of the A+B design with A=B=3 and C=D=E=1.

Web Appendix E: Prior effective sample size approximation

Prior effective sample size (ESS) is a value to quantify the information provided by a prior defined as the sample size of a previous trial where the posterior obtained is the prior distribution (Yuan, Nguyen, and Thall, 2017). Morita, Thall, and Müller (2008) have argued that a beta(a, b) distribution has effective sample size $a + b$. They proposed a definition for prior ESS of a parametric prior distribution, and compared it to a simple method (the crude method). The idea of the crude method is to define one or more probabilities, and match the prior mean and variance with those of a beta(a, b) to obtain the prior ESS. We will develop the crude method to approximate prior ESS for the logistic-DRtox model and the hierarchical-DRtox.

Let $\left((\beta_0^{(1)}, \beta_1^{(1)}), \dots, (\beta_0^{(N)}, \beta_1^{(N)}) \right)$ and $\left((\mu_z^{(1)}, \tau_z^{(1)}), \dots, (\mu_z^{(N)}, \tau_z^{(N)}) \right)$ be N -samples, where N is large, from the prior distributions for the logistic-DRtox and the hierarchical-DRtox. Let $p_{1,k}^{(n)}$ and $p_{2,k}^{(n)}$, for $n \in \{1, \dots, N\}$ and $k \in \{1, \dots, K\}$, be the probabilities for the n^{th} component of the prior vector and dose-regimen \mathbf{S}_k defined as:

$$\begin{cases} p_{1,k}^{(n)} = \pi_1 \left\{ (\beta_0^{(n)}, \beta_1^{(n)}), \bar{r}_k^M \right\} \\ p_{2,k}^{(n)} = \pi_2 \left\{ (\mu^{(n)}, \tau^{(n)}), \bar{r}_k^M \right\} \end{cases}$$

Let $\mu_{1,k}$ and $\sigma_{1,k}^2$ be the mean and variance of $(p_{1,k}^{(1)}, \dots, p_{1,k}^{(N)})$ and $\mu_{2,k}$ and $\sigma_{2,k}^2$ be the mean and

variance of $(p_{2,k}^{(1)}, \dots, p_{2,k}^{(N)})$.

The prior effective sample size for \mathbf{S}_k can be approximated as:

$$\left\{ \begin{array}{l} \text{ESS}_{1,k} = \frac{\mu_{1,k}(1 - \mu_{1,k})}{\sigma_{1,k}^2} - 1 \\ \text{ESS}_{2,k} = \frac{\mu_{2,k}(1 - \mu_{2,k})}{\sigma_{2,k}^2} - 1 \end{array} \right.$$

The global prior ESS over all dose-regimens of a scenario can be approximated as:

$$\left\{ \begin{array}{l} \text{ESS}_1 = \frac{1}{K} \sum_{k=1}^K \text{ESS}_{1,k} \\ \text{ESS}_2 = \frac{1}{K} \sum_{k=1}^K \text{ESS}_{2,k} \end{array} \right.$$

Values given in the simulation settings correspond to mean global ESS over the 6 scenarios.

Web Appendix F: Additional results


In this part we displayed additional simulation results obtained on the scenarios presented in the manuscript, results obtained on additional scenarios and results obtained with a stronger prior and those obtained when increasing the variability in toxicity.

To evaluate the impact of the prior distributions, we compared the main results with those obtained with a stronger prior distribution measured with an approximate ESS of 9, which is high for a trial including 30 patients. As the prior distributions are based on Scenario 1, stronger prior information increased the performance in the scenarios in which the dose-regimen toxicity relationship is similar to that in Scenario 1 (Scenarios 1, 3, 5 and 6), but the performance was decreased in Scenario 4. Therefore, defining prior distributions using reliable data from previous studies can increase the performance, but attention should be paid to the quality and quantity of the information used to avoid decreasing the performance.


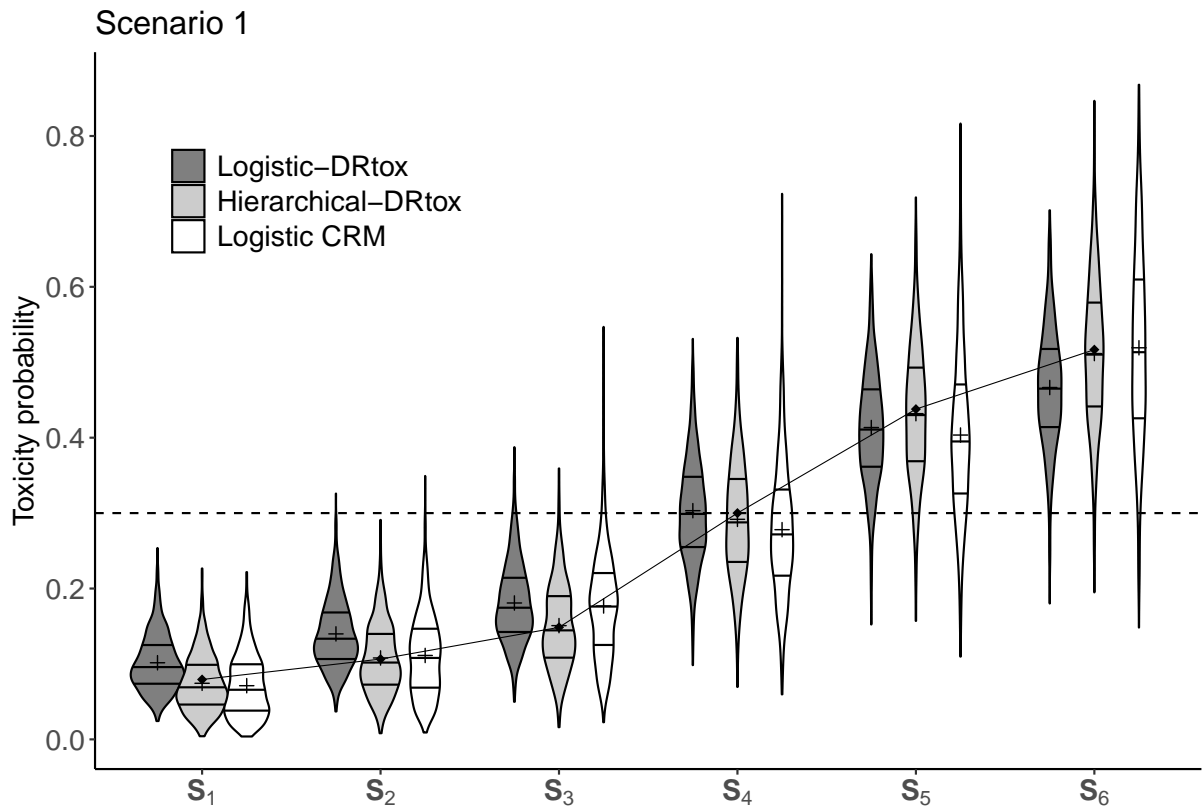
We also observed that our methods were robust to an increase in the variability in toxicity by increasing ω_α from 0.25 to 0.5 while maintaining the PK/PD variability unchanged.

All these results are displayed in the following sections



Web Appendix F.1: Additional results on the main scenarios ($\omega_\alpha = 0.25$)



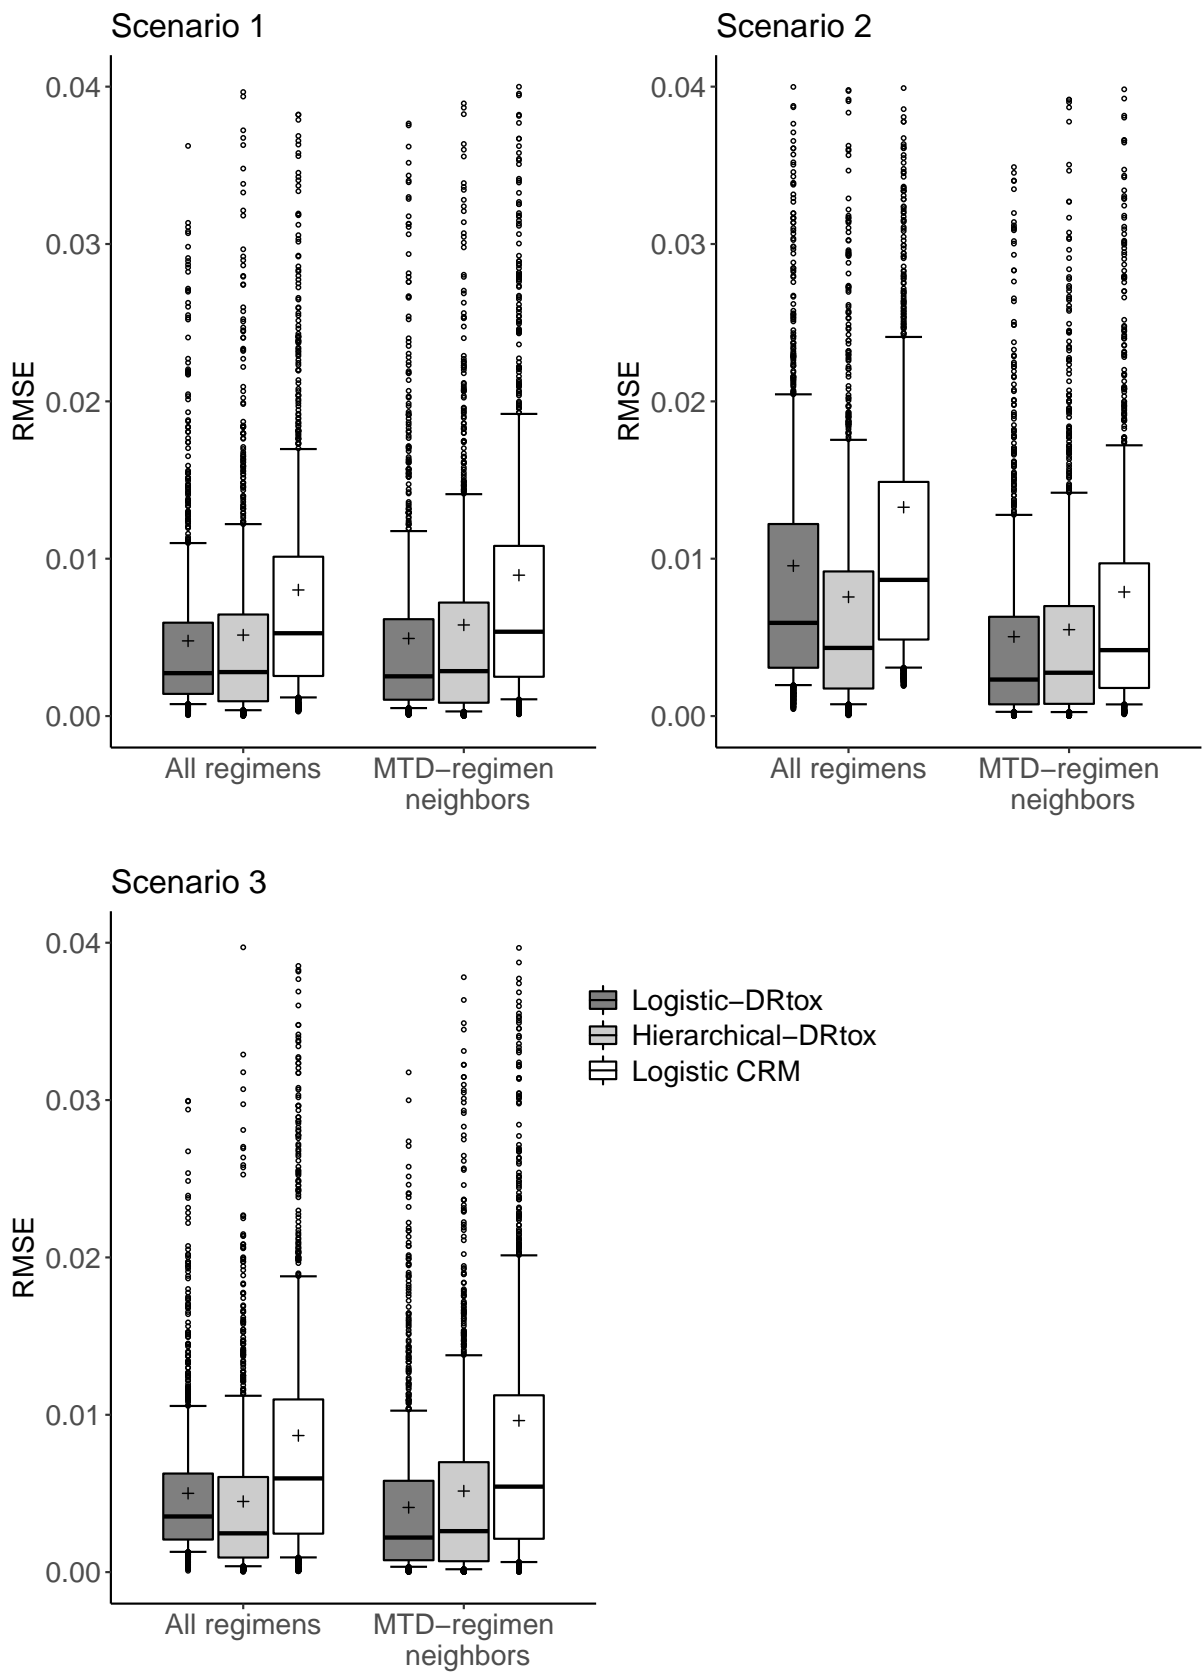
Web Figure 2: Violin plots of the estimated toxicity probabilities for an additional scenario where the dose-regimen panel missed the true MTD-regimen and scenario 3 on 1000 trials implemented with the 3+3. The predicted toxicity probability on a new regimen S_{new} is framed in dotted line. Horizontal lines on the density estimates represent the median and first and third quantiles of the distributions and the plus sign represents the mean. The dashed line represents the toxicity target and the solid line represents the true toxicity probabilities.




Web Figure 3: Violin plots of the estimated toxicity probabilities for scenarios 1 and 2 on 1000 trials implemented with the CRM including 30 patients. Horizontal lines on the density estimates represent the median and first and third quartiles of the distributions and the plus sign represents the mean. The dashed line represents the toxicity target and the solid line represents the true toxicity probabilities.



Web Figure 4: Violin plots of the estimated toxicity probabilities for scenarios 1 and 2 on 1000 trials implemented with the 3+3. Horizontal lines on the density estimates represent the median and first and third quartiles of the distributions and the plus sign represents the mean. The dashed line represents the toxicity target and the solid line represents the true toxicity probabilities.




Web Figure 5: Boxplots of the RMSE of estimated toxicity probabilities on all the dose-regimens of the panel and on the MTD-regimen and its neighbors from the panel for scenarios 1, 2 and 3 on 1000 trials implemented with the CRM including 30 patients. The plus sign represents the mean and error bars represent the first and ninth deciles.



Web Figure 6: Boxplots of the RMSE of estimated toxicity probabilities on all the dose-regimens of the panel and on the MTD-regimen and its neighbors from the panel for scenarios 1, 2 and 3 on 1000 trials implemented with the 3+3. The plus sign represents the mean and error bars represent the first and ninth deciles.

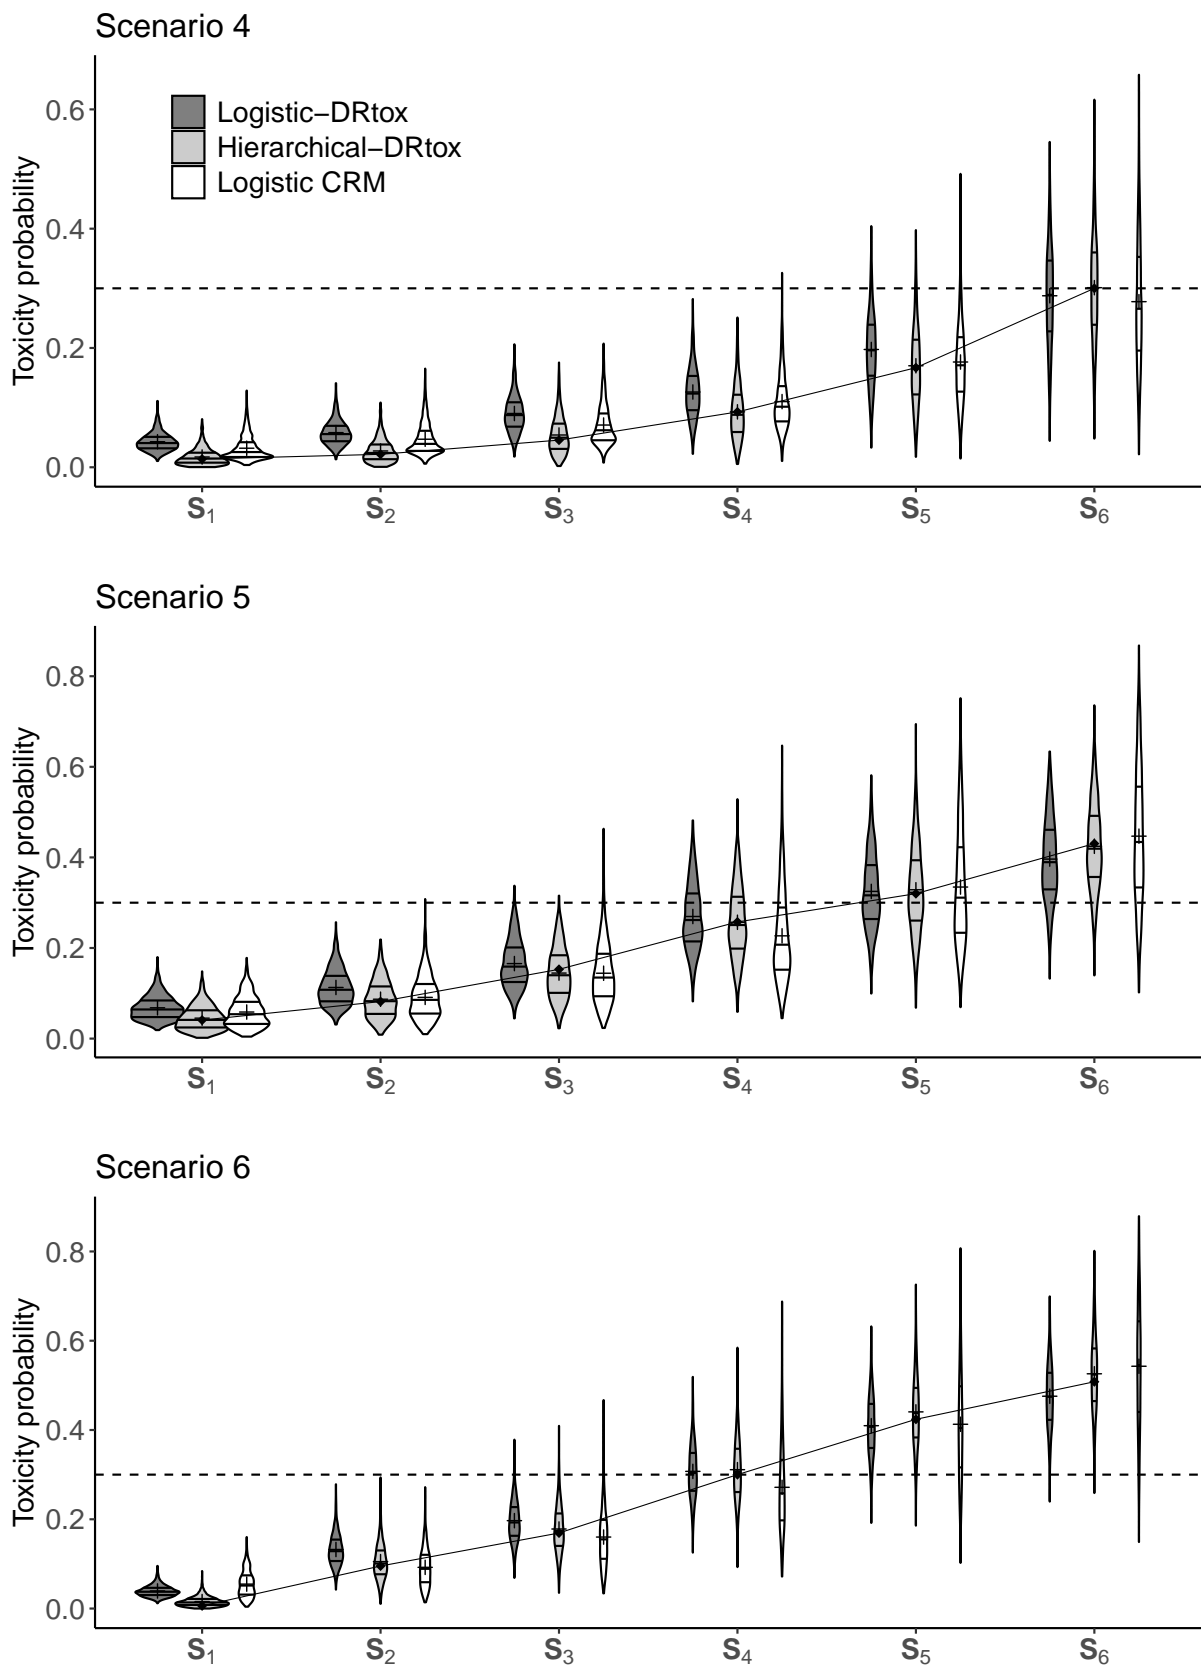
Web Appendix F.2: Additional scenarios ($\omega_\alpha = 0.25$)




Web Figure 7: The first 3 subplots represent the panel of dose-regimens for 3 additional scenarios from S_1 in spaced dashed line to S_6 in solid line, where the type of points is specific to each scenario. In the lower right corner, the dose-regimen toxicity relationship is represented for each scenario, where the MTD-regimen is the dose-regimen having the toxicity probability the closest to the target δ_T plotted in dashed line.

Web Table 2: Proportions that each dose-regimen is being selected as the MTD-regimen over the 1000 trials for the 3 additional toxicity scenarios and the 2 dose allocation designs, either the 3+3 design or the CRM. For each scenario, the PCS on the true MTD-regimen are represented in bold. For each dose allocation design, the mean sample size at each dose-regimen is displayed.


		S_1	S_2	S_3	S_4	S_5	S_6
Scenario 4		0.01	0.02	0.05	0.09	0.17	0.3
3+3	Mean sample size	3.1	3.1	3.4	3.6	3.5	3
	Logistic-DRtox	0.3	0.6	2	9.9	21.9	65.3
	Hierarchical-DRtox	0.1	0.8	2	9.3	23.8	64
	3+3	0.9	2	9.3	17.5	33.9	36.4
CRM	Mean sample size	3.1	3	3.1	3.6	5.4	11.8
	Logistic-DRtox	0	0	0	0.8	21.2	78
	Hierarchical-DRtox	0	0	0	0.4	19.9	79.7
	Logistic CRM	0	0	0	1.1	19.2	79.7
Scenario 5		0.04	0.08	0.15	0.26	0.32	0.43
3+3	Mean sample size	3.4	3.6	3.6	3.1	2	0.9
	Logistic-DRtox	3.5	5.5	21.1	30.2	23	16.7
	Hierarchical-DRtox	2.7	6.2	19.6	33.3	24	14.2
	3+3	8.4	17.2	29.2	24.1	15.4	5.7
CRM	Mean sample size	3.5	3.5	5.2	7	5.9	4.8
	Logistic-DRtox	0	0.2	11.6	34.7	30.2	23.3
	Hierarchical-DRtox	0	0.2	9.2	37.3	34	19.3
	Logistic CRM	0	0.6	9.7	28	34.4	27.3
Scenario 6		0.01	0.1	0.17	0.3	0.42	0.51
3+3	Mean sample size	3.1	3.7	3.8	3.2	1.7	0.4
	Logistic-DRtox	0.1	8.2	23.7	43	20.4	4.6
	Hierarchical-DRtox	0	8.8	25.6	47.4	16	2.2
	3+3	7.5	19.5	33.6	28.6	8.8	2
CRM	Mean sample size	3.1	3.5	5.7	8.9	6.1	2.7
	Logistic-DRtox	0	0.4	18.5	63.6	15.6	1.9
	Hierarchical-DRtox	0	0.4	17.6	67.8	13.5	0.7
	Logistic CRM	0	0.9	13.9	47.2	29.4	8.6




Web Figure 8: Violin plots of the estimated toxicity probabilities for the 3 additional scenarios on 1000 trials implemented with the CRM including 30 patients. Horizontal lines on the density estimates represent the median and first and third quartiles of the distributions and the plus sign represents the mean. The dashed line represents the toxicity target and the solid line represents the true toxicity probabilities.



Web Figure 9: Violin plots of the estimated toxicity probabilities for the 3 additional scenarios on 1000 trials implemented with the 3+3. Horizontal lines on the density estimates represent the median and first and third quartiles of the distributions and the plus sign represents the mean. The dashed line represents the toxicity target and the solid line represents the true toxicity probabilities.



Web Figure 10: Boxplots of the RMSE of estimated toxicity probabilities on all the dose-regimens of the panel and on the MTD-regimen and its neighbors from the panel for the 3 additional scenarios on 1000 trials implemented with the CRM including 30 patients. The plus sign represents the mean and error bars represent the first and ninth deciles.




Web Figure 11: Boxplots of the RMSE of estimated toxicity probabilities on all the dose-regimens of the panel and on the MTD-regimen and its neighbors from the panel for the 3 additional scenarios on 1000 trials implemented with the 3+3. The plus sign represents the mean and error bars represent the first and ninth deciles.


Web Appendix F.3: Results with ESS=9 ($\omega_\alpha = 0.25$)

Web Table 3: Proportions that each dose-regimen is being selected as the MTD-regimen, with an ESS close to 9, over the 1000 trials for all toxicity scenarios implemented with the CRM including 30 patients. For each scenario, the PCS on the true MTD-regimen are represented in bold. For each dose allocation design, the mean sample size at each dose-regimen is displayed.


		S_1	S_2	S_3	S_4	S_5	S_6
Scenario 1		0.08	0.11	0.15	0.3	0.44	0.52
CRM	Mean sample size	4.2	3.7	5.6	8.8	5.6	2.1
	Logistic-DRtox, ESS9	0	1.2	16.9	70	10.4	1.5
	Hierarchical-DRtox, ESS9	0	0.7	12.5	74.4	11.5	0.9
	Logistic CRM	0	1.4	15.1	50.4	27.1	6
Scenario 2		0.15	0.3	0.44	0.52	0.69	0.83
CRM	Mean sample size	8.7	11.1	7.5	2.3	0.3	0
	Logistic-DRtox, ESS9	9.6	70.7	18.3	1.4	0	0
	Hierarchical-DRtox, ESS9	4.1	65.4	27.4	3.1	0	0
	Logistic CRM	12.5	56	26.7	4.7	0.1	0
Scenario 3		0.07	0.11	0.2	0.3	0.42	0.56
CRM	Mean sample size	4	4	6.4	8	5.2	2.3
	Logistic-DRtox, ESS9	0.2	0.9	20	57.2	20.8	0.9
	Hierarchical-DRtox, ESS9	0.1	0.4	9.5	66	23.8	0.2
	Logistic CRM	0.1	2.3	20.3	44.5	26.4	6.4
Scenario 4		0.01	0.02	0.05	0.09	0.17	0.3
CRM	Mean sample size	3.1	3	3.1	3.6	5.4	11.8
	Logistic-DRtox, ESS9	0	0	0	2.9	41.4	55.7
	Hierarchical-DRtox, ESS9	0	0	0	0.6	56.9	42.5
	Logistic CRM	0	0	0	1.1	19.2	79.7
Scenario 5		0.04	0.08	0.15	0.26	0.32	0.43
CRM	Mean sample size	3.5	3.5	5.2	7	5.9	4.8
	Logistic-DRtox, ESS9	0	0.2	14.1	40.4	30.9	14.4
	Hierarchical-DRtox, ESS9	0	0.2	10.3	55.4	28.9	5.2
	Logistic CRM	0	0.6	9.7	28	34.4	27.3
Scenario 6		0.01	0.1	0.17	0.3	0.42	0.51
CRM	Mean sample size	3.1	3.5	5.7	8.9	6.1	2.7
	Logistic-DRtox, ESS9	0	0.2	17.8	67.7	13.1	1.2
	Hierarchical-DRtox, ESS9	0	0	9.5	83.7	6.7	0.1
	Logistic CRM	0	0.9	13.9	47.2	29.4	8.6




Web Figure 12: Violin plots of the estimated toxicity probabilities for scenarios 1, 2 and 3, with an ESS close to 9, on 1000 trials implemented with the CRM including 30 patients. The predicted toxicity probability on a new regimen S_{new} is framed in dotted line. Horizontal lines on the density estimates represent the median and first and third quantiles of the distributions and the plus sign represents the mean. The dashed line represents the toxicity target and the solid line represents the true toxicity probabilities.



Web Figure 13: Violin plots of the estimated toxicity probabilities for scenarios 4, 5 and 6, with an ESS close to 9, on 1000 trials implemented with the CRM including 30 patients. Horizontal lines on the density estimates represent the median and first and third quantiles of the distributions and the plus sign represents the mean. The dashed line represents the toxicity target and the solid line represents the true toxicity probabilities.




Web Figure 14: Boxplots of the RMSE of estimated toxicity probabilities on all the dose-regimens of the panel and on the MTD-regimen and its neighbors from the panel for scenarios 1, 2 and 3 on 1000 trials implemented with the CRM including 30 patients. The plus sign represents the mean and error bars represent the first and ninth deciles.



Web Figure 15: Boxplots of the RMSE of estimated toxicity probabilities on all the dose-regimens of the panel and on the MTD-regimen and its neighbors from the panel for scenarios 4, 5 and 6 on 1000 trials implemented with the CRM including 30 patients. The plus sign represents the mean and error bars represent the first and ninth deciles.

Web Appendix F.4: Results with $\omega_\alpha = 0.5$



Web Figure 16: Evolution of the relationship between the peak of biomarker and toxicity for ω_α varying from 0 to 0.5.

Web Table 4: Proportions that each dose-regimen is being selected as the MTD-regimen over the 1000 trials for 3 additional toxicity scenarios defined with $\omega_\alpha = 0.5$ and the 2 dose allocation designs, either the 3+3 design or the CRM. For each scenario, the PCS on the true MTD-regimen are represented in bold. For each dose allocation design, the mean sample size at each dose-regimen is displayed.

		S_1	S_2	S_3	S_4	S_5	S_6
Scenario 1 bis		0.09	0.11	0.16	0.3	0.43	0.49
3+3	Mean sample size	3.7	3.6	3.4	2.9	1.5	0.4
	Logistic-DRtox	10.5	8	19.2	37	18.6	6.7
	Hierarchical-DRtox	9.4	8.9	21.1	36.9	18.3	5.4
	3+3	16.9	16.8	31.4	24.4	8.2	2.3
CRM	Mean sample size	4.3	3.9	6	8.4	5.1	2.4
	Logistic-DRtox	0.2	1.7	18.3	56.5	16.9	6.4
	Hierarchical-DRtox	0.2	1	16.3	58.4	19.8	4.3
	Logistic CRM	0.3	1.5	16.6	48.5	23.9	9.2
Scenario 2 bis		0.16	0.3	0.43	0.49	0.6	0.73
3+3	Mean sample size	4	3.6	1.8	0.4	0.1	0
	Logistic-DRtox	29.8	40.4	23.7	5.2	0.9	0
	Hierarchical-DRtox	33.8	37.9	22.7	4.5	1.1	0
	3+3	58.4	31.3	7.9	2.4	0	0
CRM	Mean sample size	9.3	10.7	7.1	2.5	0.5	0.1
	Logistic-DRtox	18	60.5	18	2.9	0.6	0
	Hierarchical-DRtox	15.2	61.1	19.7	3.5	0.5	0
	Logistic CRM	15.9	50.3	27.2	5.4	1.2	0
Scenario 3 bis		0.09	0.13	0.2	0.3	0.41	0.54
3+3	Mean sample size	3.7	3.7	3.3	2.4	1.4	0.4
	Logistic-DRtox	11.1	10.4	24.9	28.4	20.3	4.9
	Hierarchical-DRtox	8.6	13	26	29.3	19.9	3.2
	3+3	20.2	24.3	24.1	22.8	7.9	0.7
CRM	Mean sample size	4.6	4.3	6.7	7.6	4.7	2
	Logistic-DRtox	0.6	2.8	19.3	49.3	24.5	3.5
	Hierarchical-DRtox	0.4	2.8	18.6	49.9	25.5	2.8
	Logistic CRM	0.4	4.2	20.9	44.6	22.9	7

Web Table 5: Proportions that each dose-regimen is being selected as the MTD-regimen over the 1000 trials for 3 additional toxicity scenarios defined with $\omega_\alpha = 0.5$ and the 2 dose allocation designs, either the 3+3 design or the CRM. For each scenario, the PCS on the true MTD-regimen are represented in bold. For each dose allocation design, the mean sample size at each dose-regimen is displayed.

		S_1	S_2	S_3	S_4	S_5	S_6
Scenario 4 bis		0.02	0.03	0.05	0.1	0.17	0.3
3+3	Mean sample size	3.2	3.2	3.4	3.6	3.6	3
	Logistic-DRtox	0.4	0.6	2.3	9.7	22.2	64.8
	Hierarchical-DRtox	0.2	0.8	2.2	10.1	24.8	61.9
	3+3	1	2.1	9.5	17.2	33.9	36.3
CRM	Mean sample size	3.2	3	3.1	3.6	5.4	11.6
	Logistic-DRtox	0	0	0	1.1	21.7	77.2
	Hierarchical-DRtox	0	0	0	0.9	24.4	74.7
	Logistic CRM	0	0	0	1.6	19.6	78.8
Scenario 5 bis		0.05	0.09	0.15	0.26	0.32	0.41
3+3	Mean sample size	3.4	3.6	3.6	3.1	2	0.9
	Logistic-DRtox	3.9	6.6	21.7	28	20.9	18.9
	Hierarchical-DRtox	3	7.9	21.3	30	21.7	16.1
	3+3	9.9	17.5	27.2	23.1	15.5	6.8
CRM	Mean sample size	3.6	3.5	5.2	7	5.9	4.9
	Logistic-DRtox	0	0.6	11.6	35.2	27.9	24.7
	Hierarchical-DRtox	0	0.3	10.5	36.7	29.7	22.8
	Logistic CRM	0	0.8	8.5	30	34.5	26.2
Scenario 6 bis		0.01	0.11	0.18	0.3	0.42	0.5
3+3	Mean sample size	3.1	3.8	3.7	2.9	1.7	0.4
	Logistic-DRtox	0.1	10.5	26.4	38.7	19.1	5.2
	Hierarchical-DRtox	0.1	11.6	28.4	40.8	16	3.1
	3+3	9.6	21.7	30.4	27.6	9	1.7
CRM	Mean sample size	3.2	3.5	6	8.8	5.7	2.7
	Logistic-DRtox	0	0.6	20.1	57.7	17.8	3.8
	Hierarchical-DRtox	0	0.9	19.4	61.6	16.1	2
	Logistic CRM	0	1	15.1	47	26.9	10

Web Appendix G: Computation of the hierarchical-DRtox

The Bayesian hierarchical model is defined as:

$$\begin{cases} Y_{i,j} = \begin{cases} 0 & \text{if } Z_i > \log\left(\frac{r_{i,j}}{\bar{r}_{k_{50}}^M}\right) \\ 1 & \text{if } Z_i \leq \log\left(\frac{r_{i,j}}{\bar{r}_{k_{50}}^M}\right) \end{cases} \\ Z_i \sim \mathcal{N}(\mu_z, \tau_z^2) \end{cases}$$

Therefore, the conditional distribution of Z_i is a truncated normal distribution with

$$Z_i \in \left] \max_j \left\{ \log\left(\frac{r_{i,j}}{\bar{r}_{k_{50}}^M}\right) \middle| Y_{i,j} = 0 \right\}, \min_j \left\{ \log\left(\frac{r_{i,j}}{\bar{r}_{k_{50}}^M}\right) \middle| Y_{i,j} = 1 \right\} \right[$$

In fact, under our assumptions, we only have 3 cases:

- Case 1: Patient i never experiences toxicity ($Y_i = 0$), therefore

$$Z_i > \max_{j \leq j_i} \left\{ \log\left(\frac{r_{i,j}}{\bar{r}_{k_{50}}^M}\right) \middle| Y_{i,j} = 0 \right\}$$

- Case 2: Patient i experiences toxicity at the first administration ($Y_{i,1} = 1$), therefore

$$Z_i \leq \log \left(\frac{r_{i,1}}{\bar{r}_{k_{50}}^M} \right)$$

- Case 3: Patient i experiences toxicity after the first administration ($Y_{i,0} = 1$ & $Y_i = 1$), therefore

$$\begin{cases} Z_i > \max_{j < j_i} \left\{ \log \left(\frac{r_{i,j}}{\bar{r}_{k_{50}}^M} \right) \middle| Y_{i,j} = 0 \right\} \\ Z_i \leq \log \left(\frac{r_{i,j_i}}{\bar{r}_{k_{50}}^M} \right) \end{cases}$$

To compute the model in Stan, let $Z_i^{(\text{raw})}$ be an unconstrained variable from which we will derive Z_i . We also need to introduce the bounds, let b_i^{low} and b_i^{upp} be respectively the lower and upper bound for Z_i , they are calculated from the data.

We can then compute Z_i for each case:

- Case 1: Patient i never experiences toxicity ($Y_i = 0$), therefore

$$\begin{cases} b_i^{\text{low}} = \max_{j \leq j_i} \left\{ \log \left(\frac{r_{i,j}}{\bar{r}_{k_{50}}^M} \right) \right\} \\ b_i^{\text{upp}} = +\infty \end{cases}$$

In this case, we only have $Z_i > b_i^{\text{low}}$, therefore

$$Z_i = b_i^{\text{low}} + e^{Z_i^{(\text{raw})}}$$

- Case 2: Patient i experiences toxicity at the first administration ($Y_{i,1} = 1$), therefore

$$\begin{cases} b_i^{\text{low}} = -\infty \\ b_i^{\text{upp}} = \log \left(\frac{r_{i,1}}{\bar{r}_{k_{50}}^M} \right) \end{cases}$$

In this case, we only have $Z_i < b_i^{\text{upp}}$, therefore

$$Z_i = b_i^{\text{upp}} - e^{Z_i^{(\text{raw})}}$$

- Case 3: Patient i experiences toxicity after the first administration ($Y_{i,0} = 1$ & $Y_i = 1$), therefore

$$\begin{cases} b_i^{\text{low}} = \max_{j < j_i} \left\{ \log \left(\frac{r_{i,j}}{\bar{r}_{k_{50}}^M} \right) \middle| Y_{i,j} = 0 \right\} \\ b_i^{\text{upp}} = \log \left(\frac{r_{i,j_i}}{\bar{r}_{k_{50}}^M} \right) \end{cases}$$

In this case, we have $Z_i > b_i^{\text{low}}$ & $Z_i < b_i^{\text{upp}}$, therefore

$$Z_i = b_i^{\text{low}} + \left(b_i^{\text{upp}} - b_i^{\text{low}} \right) \text{logit}^{-1} \left(Z_i^{(\text{raw})} \right)$$

Therefore, to compute the Bayesian Hierarchical model, a change of the variable $Z_i^{(\text{raw})}$ is used. In Stan, according to the reference manual, the change of variables can be applied in the sampling statement. To adjust for the curvature, the log probability accumulator is incremented with the log absolute derivative of the transform. If we note $Z_i^{(\text{raw})} = f(Z_i)$, we have to increment the log probability accumulator with

$$\log \left\{ \left| \frac{d}{dZ_i^{(\text{raw})}} f^{-1} \left(Z_i^{(\text{raw})} \right) \right| \right\}$$

Therefore, for each case we have to increment the log probability accumulator with:

- Case 1: Patient i never experiences toxicity ($Y_i = 0$)

$$\begin{aligned} \log \left\{ \left| \frac{d}{dZ_i^{(\text{raw})}} f^{-1} \left(Z_i^{(\text{raw})} \right) \right| \right\} &= \log \left\{ \left| \frac{d}{dZ_i^{(\text{raw})}} \left(b_i^{\text{low}} + e^{Z_i^{(\text{raw})}} \right) \right| \right\} \\ &= \log \left\{ \left| \frac{d}{dZ_i^{(\text{raw})}} \left(b_i^{\text{low}} + e^{Z_i^{(\text{raw})}} \right) \right| \right\} \end{aligned}$$

- Case 2: Patient i experiences toxicity at the first administration ($Y_{i,1} = 1$)

$$\begin{aligned} \log \left\{ \left| \frac{d}{dZ_i^{(\text{raw})}} f^{-1} \left(Z_i^{(\text{raw})} \right) \right| \right\} &= \log \left\{ \left| \frac{d}{dZ_i^{(\text{raw})}} \left(b_i^{\text{upp}} - e^{Z_i^{(\text{raw})}} \right) \right| \right\} \\ &= \log \left\{ \left| \frac{d}{dZ_i^{(\text{raw})}} \left(b_i^{\text{upp}} - e^{Z_i^{(\text{raw})}} \right) \right| \right\} \end{aligned}$$

- Case 3: Patient i experiences toxicity after the first administration ($Y_{i,0} = 1$ & $Y_i = 1$)

$$\begin{aligned} \log \left\{ \left| \frac{d}{dZ_i^{(\text{raw})}} f^{-1} \left(Z_i^{(\text{raw})} \right) \right| \right\} &= \log \left[\left| \frac{d}{dZ_i^{(\text{raw})}} \left\{ b_i^{\text{low}} + \left(b_i^{\text{upp}} - b_i^{\text{low}} \right) \text{logit}^{-1} \left(Z_i^{(\text{raw})} \right) \right\} \right| \right] \\ &= \log \left[\left| \left\{ b_i^{\text{upp}} - b_i^{\text{low}} \right\} \text{logit}^{-1} \left\{ Z_i^{(\text{raw})} \right\} \left\{ 1 - \text{logit}^{-1} \left(Z_i^{(\text{raw})} \right) \right\} \right| \right] \\ &= \log \left(b_i^{\text{upp}} - b_i^{\text{low}} \right) + \log \left\{ \text{logit}^{-1} \left(Z_i^{(\text{raw})} \right) \right\} \\ &\quad + \log \left\{ 1 - \text{logit}^{-1} \left(Z_i^{(\text{raw})} \right) \right\} \end{aligned}$$

Appendix D

Supporting information: Bayesian modeling of a bivariate toxicity outcome for early phase oncology trials evaluating dose regimens

Supporting information for Bayesian modeling of a bivariate toxicity outcome for early phase oncology trials evaluating dose regimens by Emma Gerard, Sarah Zohar, Christelle Lorenzato, Moreno Ursino and Marie-Karelle Riviere

Contents

Web Appendix A: Additional results	2
Web Appendix A.1: Example from a single simulated trial	2
Web Appendix A.2: Estimation of the toxicity curves	3
Web Appendix A.3: Various associations between the CRS and the DLT_o	7
Web Appendix A.4: Simpler model on the DLT_o	10
Web Appendix B: Sensitivity analysis	12
Web Appendix B.1: Sensitivity to prior effective sample size	12
Web Appendix B.2: Sensitivity to prior distribution	15
Web Appendix B.3: Sensitivity to the copula distribution	17
Web Appendix B.4: Sensitivity to the dose escalation design	20
Web Appendix C: Alternative set of dose regimens	22

Web Appendix A: Additional results

Web Appendix A.1: Example from a single simulated trial

Figure 1 shows the fit of the PD profile of patient 10 who receives dose regimen S_3 . For this patient, the global peak of cytokine is reached after administration 4, its estimated value is 600.26 pg/mL.




Figure 1: Estimated cytokine profile of patient 10 receiving S_3 and having ($Cl=1.99$, $V=3.95$, $E_{max}=645257$, $EC_{50}=10000$, $H=0.96$, $I_{max}=1$, $IC_{50}=18200$, $k_{deg}=0.21$, $K=2.43$) as individual PK/PD parameters. The dots represent the sampled cytokine responses and the continuous line shows the fitted cytokine response.

The posterior distributions of the probabilities of CRS, DLT_o , $DLT_o|noCRS$ and DLT of dose regimen S_4 are represented in Figure 2.

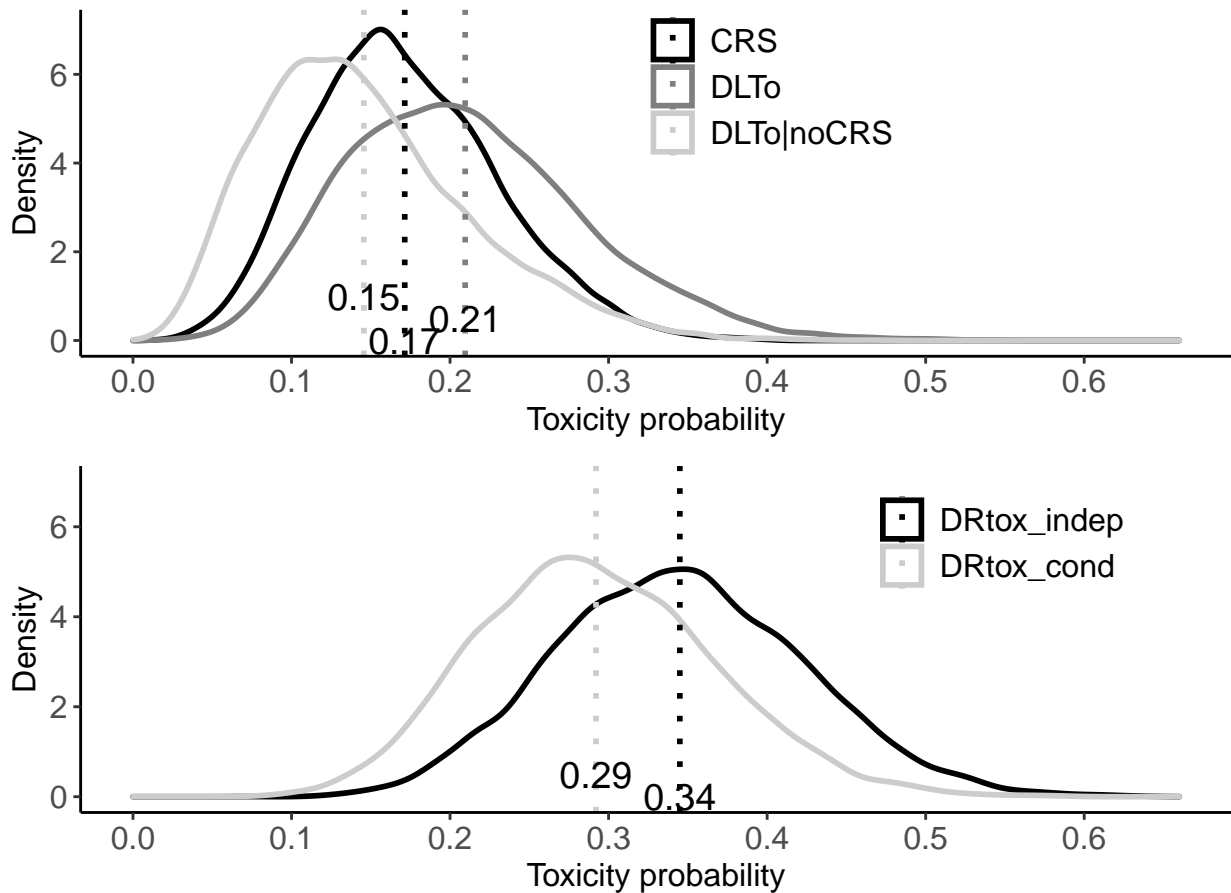


Figure 2: Estimated posterior distributions of CRS, DLT₀ and DLT₀|noCRS in the upper part of the Figure, and of DLT for the DRtox_indep and DRtox_cond in the lower part, for dose regimen S_4 . The dotted vertical lines represent the posterior means.

Web Appendix A.2: Estimation of the toxicity curves

In the main paper, we presented the results of our proposed methods in terms of the proportion of correct selection (PCS). We illustrate here the results in terms of estimation of the different probabilities of toxicity. The estimated probabilities of DLT, CRS and DLT₀ for Scenarios 2, 5 and 6 are displayed in Figures 3, 4 and 5. All three joint methods and the CRM well estimate the probability of DLT of the MTD-regimen in all scenarios. The probability of CRS of all dose regimens is well estimated by the three joint approaches via the logistic-DRtox. In Scenarios 2 and 6, both the DRtox_indep and DRtox_copula under-estimate the marginal probability of DLT₀ as they estimate it to be similar to the conditional probability of DLT₀ given no CRS. However, the DRtox_cond has a correct estimation of the conditional probability of DLT₀ given no CRS of the MTD-regimen. The under-estimation of the marginal probability of DLT₀ is due to the fact that the drug administration is stopped in case a DLT occurs (either a CRS or DLT₀) and that the CRS has a tendency to occur at the beginning of the regimens while the DLT₀ occurs at the end. Therefore, when a patient experiences a CRS, s/he does not receive the remaining administrations planned of the regimen that may have caused a DLT₀. The conditional probability of DLT₀ given that a CRS occurred can then only be estimated when a CRS and DLT₀ occur at the same time, which is rare.

In Scenario 5, where the MTD-regimen is the last one of the set, the probability of DLT of the MTD-regimen is well estimated by the three joint approaches and the CRM, but the joint approaches over-estimate the probability of DLT of S_5 , which results in a loss of PCS. Both the probabilities of CRS and DLT₀ of S_5 are over-estimated as in this scenario very few DLT are observed, therefore our joint approaches have difficulty in distinguishing regimens S_5 and S_6 .




Figure 3: Violin plots of the estimated probabilities of DLT, CRS and DLT_0 in Scenario 2 for 1000 simulated trials. All three joint approaches and the CRM estimate the probability of DLT in the first part of the figure, where the dashed line represents the toxicity target and the solid line represents the true DLT probabilities. Our three joint approaches estimate the probability of CRS with the logistic DRtox in the second part of the figure, where the solid line represents the true CRS probabilities. In the last part of the figure, both the DRtox_indep and DRtox_copula estimate the marginal probability of DLT_0 while the DRtox_cond estimates the conditional probability of DLT_0 given no CRS has occurred. The solid line represents the true marginal probabilities of DLT_0 while the dotted line represents the true conditional probabilities of DLT_0 given no CRS. In all the figure, horizontal lines on the density estimates represent the median and first and third quantiles of the distributions, and the plus sign represents the mean.




Figure 4: Violin plots of the estimated probabilities of DLT, CRS and DLT₀ in Scenario 5 for 1000 simulated trials. All three joint approaches and the CRM estimate the probability of DLT in the first part of the figure, where the dashed line represents the toxicity target and the solid line represents the true DLT probabilities. Our three joint approaches estimate the probability of CRS with the logistic DRtox in the second part of the figure, where the solid line represents the true CRS probabilities. In the last part of the figure, both the DRtox_indep and DRtox_copula estimate the marginal probability of DLT₀ while the DRtox_cond estimates the conditional probability of DLT₀ given no CRS has occurred. The solid line represents the true marginal probabilities of DLT₀ while the dotted line represents the true conditional probabilities of DLT₀ given no CRS. In all the figure, horizontal lines on the density estimates represent the median and first and third quantiles of the distributions, and the plus sign represents the mean.




Figure 5: Violin plots of the estimated probabilities of DLT, CRS and DLT_0 in Scenario 6 for 1000 simulated trials. All three joint approaches and the CRM estimate the probability of DLT in the first part of the figure, where the dashed line represents the toxicity target and the solid line represents the true DLT probabilities. Our three joint approaches estimate the probability of CRS with the logistic DRtox in the second part of the figure, where the solid line represents the true CRS probabilities. In the last part of the figure, both the DRtox_indep and DRtox_copula estimate the marginal probability of DLT_0 while the DRtox_cond estimates the conditional probability of DLT_0 given no CRS has occurred. The solid line represents the true marginal probabilities of DLT_0 while the dotted line represents the true conditional probabilities of DLT_0 given no CRS. In all the figure, horizontal lines on the density estimates represent the median and first and third quantiles of the distributions, and the plus sign represents the mean.

Web Appendix A.3: Various associations between the CRS and the DLT_o

We studied the effect of varying the association between the CRS and the DLT_o, measured by the mean risk ratio (RR), and defined three additional scenarios on Set A:

- Scenario 7: moderate positive association (RR=3.31)
- Scenario 8: independence between toxicities (RR=1)
- Scenario 9: negative association (RR=0.52)

The PCS results on these additional scenarios for our three joint approaches and the CRM are displayed in Table 1. Our approaches still outperform the CRM, and the three approaches have similar results except when increasing the correlation between the CRS and the DLT_o: the DRtox_copula and DRtox_cond have higher PCS as they account for the association between toxicities. We can also note that the DRtox_copula, that assumes a positive association between toxicities, still has good results on Scenario 9 where there is a negative association between the CRS and the DLT_o.

The estimations of the DLT probabilities in case of independence (Scenario 8), small association (Scenario 1) and high association (Scenario 2) are represented in Figure 6. The estimations on the six dose regimens and the predictions on \mathbf{S}_{new1} and \mathbf{S}_{new2} are shown. The root-mean square error (RMSE) of the estimated probabilities on $\mathbf{S}_3, \mathbf{S}_4, \mathbf{S}_5$ (neighbors) and on $\mathbf{S}_4, \mathbf{S}_{\text{new1}}, \mathbf{S}_{\text{new2}}$ (predict) are represented in Figure 7. We can observe that all methods, DRtox_indep, DRtox_copula and DRtox_cond, have good estimations around the MTD-regimen in case of various associations between the CRS and the DLT_o.

Table 1: Proportions of selecting each dose regimen as the MTD-regimen over the 1000 trials in three additional toxicity scenarios with various associations between the CRS and the DLT_o. For each scenario, the marginal probabilities of DLT, CRS and DLT_o are defined, and the association between the CRS and DLT_o is represented by the average risk ratio (RR). Results are presented for the 3 joint approaches (DRtox_indep, DRtox_copula and DRtox_cond) and the CRM. The proportions of correct selection (PCS) of the MTD-regimen are represented in bold.

Scenario	Set	RR	Method		\mathbf{S}_1	\mathbf{S}_2	\mathbf{S}_3	\mathbf{S}_4	\mathbf{S}_5	\mathbf{S}_6
7	A	3.31		p_T	0.10	0.14	0.18	0.30	0.44	0.57
				$p_T^{(1)}$	0.06	0.08	0.11	0.19	0.24	0.32
				$p_T^{(2)}$	0.05	0.07	0.10	0.19	0.37	0.53
			DRtox_indep	0	4	25	54	15	3	
			DRtox_copula	0	1	18	56	21	4	
			DRtox_cond	0	3	23	55	17	3	
			Logistic CRM	0	4	20	46	23	7	
8	A	1.00		p_T	0.10	0.14	0.18	0.30	0.45	0.59
				$p_T^{(1)}$	0.04	0.07	0.09	0.16	0.20	0.27
				$p_T^{(2)}$	0.06	0.08	0.10	0.16	0.31	0.44
			DRtox_indep	0	3	24	57	15	1	
			DRtox_copula	0	1	16	58	22	2	
			DRtox_cond	0	3	23	57	15	1	
			Logistic CRM	0	3	22	47	23	5	
9	A	0.52		p_T	0.10	0.14	0.18	0.30	0.45	0.59
				$p_T^{(1)}$	0.04	0.06	0.09	0.16	0.20	0.27
				$p_T^{(2)}$	0.06	0.08	0.09	0.15	0.28	0.39
			DRtox_indep	0	3	21	58	16	1	
			DRtox_copula	0	1	15	58	23	3	
			DRtox_cond	0	4	22	58	16	1	
			Logistic CRM	0	4	20	48	23	5	




Figure 6: Violin plots of the estimated probabilities of DLT when increasing association between the CRS and DLT_o (Scenarios 8, 1 and 2) for the six dose regimens of the panel and two additional dose regimens (S_{new1} and S_{new2}), on 1000 trials with the three proposed joint approaches and the CRM. The predicted DLT probabilities of the new dose regimens are framed in dotted line. Horizontal lines on the density estimates represent the median and first and third quantiles of the distributions, and the plus sign represents the mean. The dashed line represents the toxicity target, and the solid line represents the true DLT probabilities




Figure 7: Boxplots of the RMSE of the estimated DLT probabilities on S_3, S_4, S_5 (neighbors) and on S_4, S_{new1}, S_{new2} (predict) when increasing association between the CRS and DLT_o (Scenarios 8, 1 and 2) on 1000 trials. The plus sign represents the mean and error bars represent the first and ninth deciles.

Web Appendix A.4: Simpler model on the DLT_o

To compute the DRtox_cond approach, we modeled the DLT_o with a cumulative model using the cumulative dose to account for the dose regimen. We evaluated the effect of a simpler model on the DLT_o without taking into account the multiple administrations. We defined the conditional probability of DLT_o given that no CRS has occurred as follows:

$$p_{i\star}^{(2)} = \mathbb{P}\left(Y_i^{(2)} = 1 \mid Y_i^{(1)} = 0\right) \quad (1)$$

We then defined the following model on the conditional probability of DLT_o that is very similar to the 2-parameter logistic model of the CRM:

$$\text{logit}\left(p_{i\star}^{(2)}\right) = a + b \text{logit}\left(\pi_{k_i}^{(2)}\right) \quad (2)$$

where $\pi_{k_i}^{(2)}$ is the prior guess of the DLT_o probability of dose regimen S_{k_i} that is planned for patient i . We initially assume that the probabilities of CRS and DLT_o are independent and equal, therefore $\pi_{k_i}^{(2)} = 1 - \sqrt{1 - \pi_{k_i}}$. For the prior distributions, we considered $a \sim \mathcal{N}(0, \sqrt{10})$ and $b \sim \gamma(1, 1)$ to ensure positivity.

Let DRtox_cond_simple be the joint approach built on the conditional formulation and using the simpler model on the DLT_o defined in Equation 2. The PCS of the DRtox_cond_simple and DRtox_cond are displayed in Table 2 for the six main scenarios. The only case where the simpler model is better is in Scenario 5 where the MTD-regimen is the last one of the set and therefore few DLT_o are observed. In this case, distinguishing the different regimens becomes challenging for the cumulative model in the DRtox_cond.

Table 2: Proportions of selecting each dose regimen as the MTD-regimen over the 1000 trials in the six main toxicity scenarios. For each scenario, the marginal probabilities of DLT, CRS and DLT_o are defined, and the association between the CRS and DLT_o is represented by the average risk ratio (RR). Results are presented for the joint approach defined from the conditional formulation using either the cumulative model (DRtox_cond) or the simpler model on the DLT_o (DRtox_cond_simple). The proportions of correct selection (PCS) of the MTD-regimen are represented in bold.

Scenario	Set	RR	Method	S_1	S_2	S_3	S_4	S_5	S_6
1	A	1.85	p_T	0.10	0.14	0.18	0.30	0.45	0.60
			$p_T^{(1)}$	0.05	0.07	0.10	0.18	0.22	0.30
			$p_T^{(2)}$	0.06	0.08	0.10	0.17	0.34	0.50
			DRtox_cond	0	3	25	55	16	2
			DRtox_cond_simple	0	4	27	50	16	2
2	A	5.91	p_T	0.10	0.13	0.18	0.30	0.42	0.53
			$p_T^{(1)}$	0.07	0.10	0.13	0.22	0.28	0.36
			$p_T^{(2)}$	0.04	0.07	0.10	0.21	0.39	0.52
			DRtox_cond	0	3	23	52	18	4
			DRtox_cond_simple	1	4	23	49	20	4
3	A	1.81	p_T	0.11	0.15	0.18	0.30	0.45	0.59
			$p_T^{(1)}$	0.03	0.05	0.07	0.13	0.16	0.23
			$p_T^{(2)}$	0.08	0.11	0.13	0.22	0.39	0.54
			DRtox_cond	1	3	22	58	15	2
			DRtox_cond_simple	1	5	27	50	15	2
4	A	1.90	p_T	0.09	0.13	0.17	0.30	0.44	0.59
			$p_T^{(1)}$	0.07	0.10	0.13	0.23	0.29	0.37
			$p_T^{(2)}$	0.03	0.04	0.06	0.12	0.27	0.43
			DRtox_cond	0	3	24	55	15	2
			DRtox_cond_simple	0	4	25	52	16	3
5	A	1.97	p_T	0.03	0.04	0.05	0.11	0.17	0.30
			$p_T^{(1)}$	0.02	0.03	0.05	0.09	0.11	0.17
			$p_T^{(2)}$	0.00	0.01	0.01	0.02	0.08	0.18
			DRtox_cond	0	0	0	4	29	67
			DRtox_cond_simple	0	0	0	4	20	76
6	B	1.70	p_T	0.16	0.30	0.43	0.55	0.73	0.86
			$p_T^{(1)}$	0.09	0.17	0.25	0.33	0.38	0.48
			$p_T^{(2)}$	0.08	0.18	0.29	0.41	0.67	0.84
			DRtox_cond	19	65	15	1	0	0
			DRtox_cond_simple	21	63	15	1	0	0

Web Appendix B: Sensitivity analysis

Web Appendix B.1: Sensitivity to prior effective sample size

We evaluated the effect of varying the amount of information provided by the prior distributions that we measured by approximating the effective sample size (ESS). We studied three different ESS to evaluate the effect of almost no prior information (ESS=0.2), medium prior information (ESS=2) and strong prior information (ESS=7).

The case of almost no prior information, ESS=0.2, was obtained with $\sigma_{\beta_{0,1}} = 10$ $\alpha = 5$ for the CRS model and $\sigma_{\beta_{0,2}} = 10$ and $\sigma_{\beta_{1,2}} = 1$ for the DLT_o model. The case of medium prior association, ESS=2, was obtained with $\sigma_{\beta_{0,1}} = 2$ $\alpha = 5$ for the CRS model and $\sigma_{\beta_{0,2}} = 2$ and $\sigma_{\beta_{1,2}} = 1$ for the DLT_o model. The case of strong prior information, ESS=7, was obtained with $\sigma_{\beta_{0,1}} = 1$ $\alpha = 5$ for the CRS model and $\sigma_{\beta_{0,2}} = 1$ and $\sigma_{\beta_{1,2}} = 0.45$ for the DLT_o model.

PCS with these increasing ESS for the DRtox_indep, DRtox_copula and DRtox_cond can be found in Table 3 for Scenarios 1, 2 and 3 and in Table 4 for Scenarios 4, 5 and 6. Increasing the prior ESS leads to better results when the prior guesses of DLT probabilities are close to the truth (Scenarios 1-4 where \mathbf{S}_4 is the true MTD-regimen), but also when the initial guesses of DLT probabilities underestimate the true DLT probabilities (Scenario 6 where \mathbf{S}_2 is the true MTD-regimen). However, increasing the prior ESS leads to poorer results when the initial guesses of DLT probabilities overestimate the true DLT probabilities (Scenario 5 where \mathbf{S}_6 is the true MTD-regimen).

Table 3: Proportions of selecting each dose regimen as the MTD-regimen over the 1000 trials in Scenarios 1, 2 and 3 for increasing ESS. For each scenario, the marginal probabilities of DLT, CRS and DLT_o are defined, and the association between the CRS and DLT_o is represented by the average risk ratio (RR). Results are presented for the 3 joint approaches (DRtox_indep, DRtox_copula and DRtox_cond) and the CRM. The proportions of correct selection (PCS) of the MTD-regimen are represented in bold.

Scenario	Set	RR	Method	ESS		S_1	S_2	S_3	S_4	S_5	S_6
1	A	1.85			p_T	0.10	0.14	0.18	0.30	0.45	0.60
					$p_T^{(1)}$	0.05	0.07	0.10	0.18	0.22	0.30
					$p_T^{(2)}$	0.06	0.08	0.10	0.17	0.34	0.50
			DRtox_indep	0.2	0	3	22	53	18	2	
				2	0	4	24	55	15	2	
				7	0	3	30	57	9	0	
			DRtox_copula	0.2	0	2	19	51	24	4	
				2	0	2	20	55	20	3	
				7	0	1	21	61	16	1	
			DRtox_cond	0.2	0	3	22	52	19	3	
				2	0	3	25	55	15	2	
				7	0	3	31	57	9	0	
2	A	5.91			p_T	0.10	0.13	0.18	0.30	0.42	0.53
					$p_T^{(1)}$	0.07	0.10	0.13	0.22	0.28	0.36
					$p_T^{(2)}$	0.04	0.07	0.10	0.21	0.39	0.52
			DRtox_indep	0.2	1	4	23	46	20	6	
				2	0	4	27	48	16	3	
				7	0	5	31	54	9	1	
			DRtox_copula	0.2	0	3	15	50	24	8	
				2	0	3	17	52	24	5	
				7	0	2	20	60	15	3	
			DRtox_cond	0.2	1	3	20	49	22	6	
				2	0	3	24	52	18	4	
				7	0	3	30	56	9	1	
3	A	1.81			p_T	0.11	0.15	0.18	0.30	0.45	0.59
					$p_T^{(1)}$	0.03	0.05	0.07	0.13	0.16	0.23
					$p_T^{(2)}$	0.08	0.11	0.13	0.22	0.39	0.54
			DRtox_indep	0.2	1	4	20	55	18	2	
				2	1	3	22	58	15	2	
				7	1	2	26	62	8	0	
			DRtox_copula	0.2	1	2	15	53	24	4	
				2	1	2	16	59	20	2	
				7	0	1	16	64	18	1	
			DRtox_cond	0.2	1	3	19	56	19	2	
				2	1	3	21	58	16	2	
				7	1	2	26	62	9	0	

Table 4: Proportions of selecting each dose regimen as the MTD-regimen over the 1000 trials in Scenarios 4, 5 and 6 for increasing ESS. For each scenario, the marginal probabilities of DLT, CRS and DLT_o are defined, and the association between the CRS and DLT_o is represented by the average risk ratio (RR). Results are presented for the 3 joint approaches (DRtox_indep, DRtox_copula and DRtox_cond) and the CRM. The proportions of correct selection (PCS) of the MTD-regimen are represented in bold.

Scenario	Set	RR	Method	ESS	S_1	S_2	S_3	S_4	S_5	S_6
4	A	1.90		p_T	0.09	0.13	0.17	0.30	0.44	0.59
				$p_T^{(1)}$	0.07	0.10	0.13	0.23	0.29	0.37
				$p_T^{(2)}$	0.03	0.04	0.06	0.12	0.27	0.43
			DRtox_indep	0.2	1	3	22	52	19	4
				2	0	3	24	56	14	2
				7	0	3	29	59	8	0
	DRtox_copula	0.2	0	2	17	52	24	5		
		2	0	2	19	56	20	4		
		7	0	2	21	61	14	2		
	DRtox_cond	0.2	1	3	21	53	18	4		
		2	0	3	24	55	14	2		
		7	0	3	31	57	8	0		
5	A	1.97		p_T	0.03	0.04	0.05	0.11	0.17	0.30
				$p_T^{(1)}$	0.02	0.03	0.05	0.09	0.11	0.17
				$p_T^{(2)}$	0.00	0.01	0.01	0.02	0.08	0.18
			DRtox_indep	0.2	0	0	0	3	20	78
				2	0	0	0	4	29	67
				7	0	0	0	10	43	47
	DRtox_copula	0.2	0	0	0	2	15	83		
		2	0	0	0	3	20	77		
		7	0	0	0	5	32	63		
	DRtox_cond	0.2	0	0	0	3	19	78		
		2	0	0	0	4	29	67		
		7	0	0	0	10	43	47		
6	B	1.70		p_T	0.16	0.30	0.43	0.55	0.73	0.86
				$p_T^{(1)}$	0.09	0.17	0.25	0.33	0.38	0.48
				$p_T^{(2)}$	0.08	0.18	0.29	0.41	0.67	0.84
			DRtox_indep	0.2	24	60	14	1	0	0
				2	20	64	15	1	0	0
				7	13	72	15	0	0	0
	DRtox_copula	0.2	18	60	20	2	0	0		
		2	14	62	21	2	0	0		
		7	6	68	25	1	0	0		
	DRtox_cond	0.2	22	61	15	1	0	0		
		2	19	65	16	1	0	0		
		7	12	71	16	0	0	0		

Web Appendix B.2: Sensitivity to prior distribution

To evaluate the effect of the prior distribution on the DRtox_cond, we compared the results when using a gamma distribution on the slope or a normal distribution on the logarithm of the slope for both the CRS model and the DLT_o model.

For the CRS model, we considered:

- Gamma prior: $\text{logit} \left(\mathbb{P} \left(Y_i^{(1)} = 1 \right) \right) = \beta_{0,1} + \beta_{1,1} \log \left(\frac{r_i^M}{\bar{r}_{kT}^M} \right)$, where $\beta_{1,1} \sim \gamma \left(\alpha_1, \frac{\alpha_1}{\bar{\beta}_{1,1}} \right)$
- Normal prior: $\text{logit} \left(\mathbb{P} \left(Y_i^{(1)} = 1 \right) \right) = \beta_{0,1} + \exp(\beta_{1,1}) \log \left(\frac{r_i^M}{\bar{r}_{kT}^M} \right)$, where $\beta_{1,1} \sim \mathcal{N} \left(\bar{\beta}_{1,1}, \sigma_{\beta_1}^2 \right)$

For the DLT_o model, we consider:

- Gamma prior: $\text{logit} \left(p_{i,j^*}^{(2)\text{cum}} \right) = \beta_{0,2^*} + \beta_{1,2^*} \log \left(\frac{\sum_{l=1}^j d_{i,l}}{D_{kT}} \right)$, where $\beta_{1,2^*} \sim \gamma \left(\alpha_1, \frac{\alpha_1}{\bar{\beta}_{1,2^*}} \right)$
- Normal prior: $\text{logit} \left(p_{i,j^*}^{(2)\text{cum}} \right) = \beta_{0,2^*} + \exp(\beta_{1,2^*}) \log \left(\frac{\sum_{l=1}^j d_{i,l}}{D_{kT}} \right)$, where $\beta_{1,2^*} \sim \mathcal{N} \left(\bar{\beta}_{1,2^*}, \sigma_{\beta_1}^2 \right)$

For both the CRS and DLT_o models, we considered $\alpha_1 = 5$, $\sigma_{\beta_1} = 1$. For the intercept, we considered $\beta_{0,1} \sim \mathcal{N} \left(\bar{\beta}_{0,1}, \sigma_{\beta_0}^2 \right)$ and $\beta_{0,2^*} \sim \mathcal{N} \left(\bar{\beta}_{0,2^*}, \sigma_{\beta_0}^2 \right)$, where $\sigma_{\beta_0} = 2$.

PCS results with these various prior distributions can be found in Table 5 for the DRtox_cond approach on Scenarios 1-6. The prior distribution has little impact on the results for almost all scenarios. In Scenario 5, where the true MTD-regimen is the last dose regimen of the panel and therefore only few DLT are observed, choosing a gamma prior for the DLT_o model can lead to better results.

Table 5: Proportions of selecting each dose regimen as the MTD-regimen over the 1000 trials in the six main toxicity scenarios for various prior distributions (gamma or lognormal). For each scenario, the marginal probabilities of DLT, CRS and DLT_o are defined, and the association between the CRS and DLT_o is represented by the average risk ratio (RR). Results are presented for the DRtox_cond. The proportions of correct selection (PCS) at the MTD-regimen are represented in bold.

Scenario	Set	RR	Method	CRS	DLT _o	p_T	S_1	S_2	S_3	S_4	S_5	S_6	
1	A	1.85	DRtox_cond			p_T	0.10	0.14	0.18	0.30	0.45	0.60	
						$p_T^{(1)}$	0.05	0.07	0.10	0.18	0.22	0.30	
						$p_T^{(2)}$	0.06	0.08	0.10	0.17	0.34	0.50	
						γ	γ	0	3	23	53	18	2
						γ	\mathcal{N}	0	3	25	55	15	2
\mathcal{N}	γ	0	3	27	54	14	1						
\mathcal{N}	\mathcal{N}	0	3	28	55	12	1						
2	A	5.91	DRtox_cond			p_T	0.10	0.13	0.18	0.30	0.42	0.53	
						$p_T^{(1)}$	0.07	0.10	0.13	0.22	0.28	0.36	
						$p_T^{(2)}$	0.04	0.07	0.10	0.21	0.39	0.52	
						γ	γ	0	3	22	51	20	4
						γ	\mathcal{N}	0	3	24	52	18	4
\mathcal{N}	γ	0	4	26	54	14	2						
\mathcal{N}	\mathcal{N}	0	4	27	54	13	2						
3	A	1.81	DRtox_cond			p_T	0.11	0.15	0.18	0.30	0.45	0.59	
						$p_T^{(1)}$	0.03	0.05	0.07	0.13	0.16	0.23	
						$p_T^{(2)}$	0.08	0.11	0.13	0.22	0.39	0.54	
						γ	γ	1	2	19	57	18	2
						γ	\mathcal{N}	1	3	21	58	16	2
\mathcal{N}	γ	1	3	22	57	16	1						
\mathcal{N}	\mathcal{N}	1	3	23	59	13	1						
4	A	1.90	DRtox_cond			p_T	0.09	0.13	0.17	0.30	0.44	0.59	
						$p_T^{(1)}$	0.07	0.10	0.13	0.23	0.29	0.37	
						$p_T^{(2)}$	0.03	0.04	0.06	0.12	0.27	0.43	
						γ	γ	1	3	24	53	17	3
						γ	\mathcal{N}	0	3	24	55	14	2
\mathcal{N}	γ	0	4	26	56	12	2						
\mathcal{N}	\mathcal{N}	0	4	27	57	11	1						
5	A	1.97	DRtox_cond			p_T	0.03	0.04	0.05	0.11	0.17	0.30	
						$p_T^{(1)}$	0.02	0.03	0.05	0.09	0.11	0.17	
						$p_T^{(2)}$	0.00	0.01	0.01	0.02	0.08	0.18	
						γ	γ	0	0	0	4	21	75
						γ	\mathcal{N}	0	0	0	4	29	67
\mathcal{N}	γ	0	0	0	4	25	70						
\mathcal{N}	\mathcal{N}	0	0	0	5	31	64						
6	B	1.70	DRtox_cond			p_T	0.16	0.30	0.43	0.55	0.73	0.86	
						$p_T^{(1)}$	0.09	0.17	0.25	0.33	0.38	0.48	
						$p_T^{(2)}$	0.08	0.18	0.29	0.41	0.67	0.84	
						γ	γ	19	64	16	1	0	0
						γ	\mathcal{N}	19	65	16	1	0	0
\mathcal{N}	γ	19	66	15	1	0	0						
\mathcal{N}	\mathcal{N}	18	66	15	0	0	0						

Web Appendix B.3: Sensitivity to the copula distribution

For the joint modeling approach based on a copula distribution, we evaluated two copula distributions defined as follows:

- The Clayton distribution:

$$C_\alpha \left(p_k^{(1)}, p_k^{(2)} \right) = \left(\max \left(p_k^{(1)-\gamma} + p_k^{(2)-\gamma} - 1, 0 \right) \right)^{-1/\gamma} \quad (3)$$

where $\gamma > 0$ for positive association and $\gamma \in [-1, 0[$ for negative association.

- The Farlie–Gumbel–Morgenstern distribution:

$$C_\alpha \left(p_k^{(1)}, p_k^{(2)} \right) = p_k^{(1)} p_k^{(2)} + p_k^{(1)} \left(1 - p_k^{(1)} \right) p_k^{(2)} \left(1 - p_k^{(2)} \right) \frac{\exp(\psi) - 1}{\exp(\psi) + 1} \quad (4)$$

where $\psi = 0$ for independence, $\psi > 0$ for positive association, $\psi < 0$ for negative association.

We also evaluated various prior information on each distribution and defined 4 final joint modeling approaches based on the copula distribution as follows:

- DRtox_Clayton1: $\gamma \sim \gamma(0.1, 0.1)$
- DRtox_Clayton2: $\gamma \sim \gamma(1, 1)$ (defined as DRtox_copula in the main paper)
- DRtox_Gumbel1: $\psi \sim \mathcal{N}^+(0, 3^2)$
- DRtox_Gumbel2: $\psi \sim \mathcal{N}^+(0, 1)$

The PCS of these four variants compared to the DRtox_indep are displayed in Table 6. All 4 copula approaches have similar PCS, but the DRtox_Clayton1 and DRtox_Gumbel2 have results very close to the DRtox_indep.

The histogram of the estimated parameter of the Clayton copula for Scenarios 8, 1, 2 and the histogram of the estimated parameter of the Clayton copula for the two distributions are represented in Figure 8. We can observe on the upper part of the figure that increasing the association between toxicities has little impact on the estimation of the Clayton parameter. The lower part of the figure represents the estimation of the Clayton parameter when both toxicities are highly associated in Scenario 2 for two prior distributions. We can observe that the estimation of the parameter is highly impacted by the prior distribution chosen, even when both toxicities are strongly associated. The difficulty in estimating the copula parameter can be explained by the fact that the CRS and DLT_o rarely occur at the same time as the CRS has a tendency to occur at the beginning of the regimen while the DLT_o occurs at the end.

Table 6: Proportions of selecting each dose regimen as the MTD-regimen over the 1000 trials in the six main toxicity scenarios. For each scenario, the marginal probabilities of DLT, CRS and DLT_o are defined, and the association between the CRS and DLT_o is represented by the average risk ratio (RR). Results are presented for the DRtox_indep and the four variants of the DRtox_copula. The proportions of correct selection (PCS) of the MTD-regimen are represented in bold.

Scenario	Set	RR	Method	\mathcal{S}_1	\mathcal{S}_2	\mathcal{S}_3	\mathcal{S}_4	\mathcal{S}_5	\mathcal{S}_6	
1	A	1.85		p_T	0.10	0.14	0.18	0.30	0.45	0.60
				$p_T^{(1)}$	0.05	0.07	0.10	0.18	0.22	0.30
				$p_T^{(2)}$	0.06	0.08	0.10	0.17	0.34	0.50
			DRtox_indep	0	4	24	55	16	2	
			DRtox_Clayton1	0	3	24	55	17	2	
			DRtox_Clayton2	0	2	20	55	20	3	
			DRtox_Gumbel1	0	4	25	55	15	2	
DRtox_Gumbel2	0	3	26	55	14	1				
2	A	5.91		p_T	0.10	0.13	0.18	0.30	0.42	0.53
				$p_T^{(1)}$	0.07	0.10	0.13	0.22	0.28	0.36
				$p_T^{(2)}$	0.04	0.07	0.10	0.21	0.39	0.52
			DRtox_indep	0	4	27	48	16	3	
			DRtox_Clayton1	0	4	23	50	18	4	
			DRtox_Clayton2	0	3	17	52	24	5	
			DRtox_Gumbel1	0	4	28	48	16	3	
DRtox_Gumbel2	0	4	28	49	15	4				
3	A	1.81		p_T	0.11	0.15	0.18	0.30	0.45	0.59
				$p_T^{(1)}$	0.03	0.05	0.07	0.13	0.16	0.23
				$p_T^{(2)}$	0.08	0.11	0.13	0.22	0.39	0.54
			DRtox_indep	1	3	22	57	15	2	
			DRtox_Clayton1	1	3	21	57	17	2	
			DRtox_Clayton2	1	2	16	58	20	2	
			DRtox_Gumbel1	1	3	22	58	14	2	
DRtox_Gumbel2	1	4	23	57	14	1				
4	A	1.90		p_T	0.09	0.13	0.17	0.30	0.44	0.59
				$p_T^{(1)}$	0.07	0.10	0.13	0.23	0.29	0.37
				$p_T^{(2)}$	0.03	0.04	0.06	0.12	0.27	0.43
			DRtox_indep	0	3	24	56	14	2	
			DRtox_Clayton1	0	3	22	56	15	3	
			DRtox_Clayton2	0	2	19	56	20	4	
			DRtox_Gumbel1	0	3	25	55	14	2	
DRtox_Gumbel2	0	3	25	55	14	2				
5	A	1.97		p_T	0.03	0.04	0.05	0.11	0.17	0.30
				$p_T^{(1)}$	0.02	0.03	0.05	0.09	0.11	0.17
				$p_T^{(2)}$	0.00	0.01	0.01	0.02	0.08	0.18
			DRtox_indep	0	0	0	4	29	67	
			DRtox_Clayton1	0	0	0	4	27	69	
			DRtox_Clayton2	0	0	0	3	20	77	
			DRtox_Gumbel1	0	0	0	4	29	67	
DRtox_Gumbel2	0	0	0	5	30	66				
6	B	1.70		p_T	0.16	0.30	0.43	0.55	0.73	0.86
				$p_T^{(1)}$	0.09	0.17	0.25	0.33	0.38	0.48
				$p_T^{(2)}$	0.08	0.18	0.29	0.41	0.67	0.84
			DRtox_indep	20	64	15	1	0	0	
			DRtox_Clayton1	19	64	16	1	0	0	
			DRtox_Clayton2	14	63	21	2	0	0	
			DRtox_Gumbel1	18	64	14	1	0	0	
DRtox_Gumbel2	22	64	14	0	0	0				


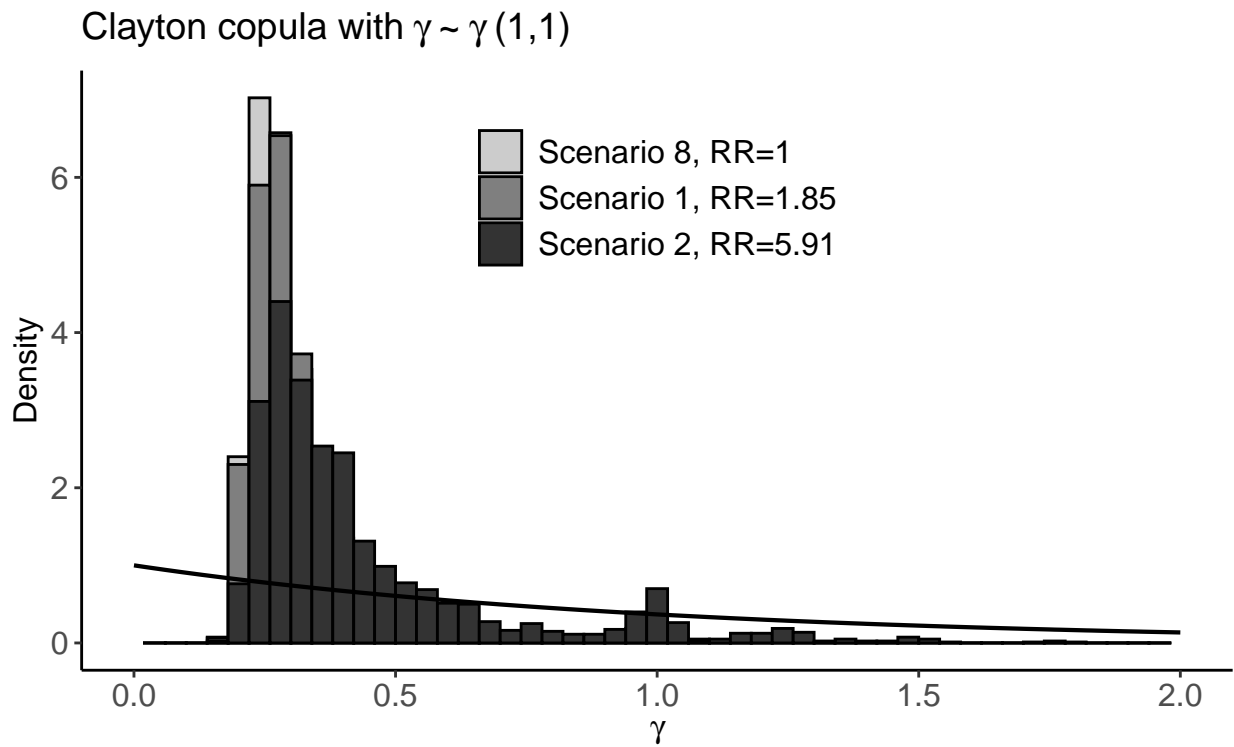


Figure 8: The upper part of the figure represents the histogram of the estimated of the Clayton distribution using $\gamma(1,1)$ for the prior distribution (named as DRtox_Clayton2) for Scenarios 8, 1 and 2. The prior distribution is represented in solid line. The lower part of the figure represents the histogram of the estimated parameter of the Clayton distribution in Scenario 2 (with a mean RR of 5.91) for the two prior distributions. Each prior distribution is represented in solid line.

Web Appendix B.4: Sensitivity to the dose escalation design

We evaluated the results obtained when the trials were simulated under an empiric CRM. The probability of DLT at dose regimen \mathbf{S}_k is defined as $p_k = \frac{\exp(\beta)}{\pi_k}$, where π_k is the initial guess of DLT probability (skeleton) and $\beta \sim \mathcal{N}(0, 1.34)$. The skeleton is the same than the one of the logistic CRM, that is (0.06,0.12,0.20,0.30,0.40,0.50).

The PCS of our proposed methods applied at the end of the empiric CRM can be found in Table 7. In all scenarios, except Scenario 5, the PCS of our proposed methods are higher than the one of the empiric CRM. In Scenario 2, the performance of the empiric CRM is similar to the DRtox_indep. The empiric CRM gives better PCS than the logistic CRM in Scenarios 1, 2 and 4, but the impact on the performance of our proposed methods is limited. In Scenario 3 and 5, the empiric and logistic CRM have the same PCS, but the performance of our proposed methods are higher after the logistic CRM in Scenario 3 and higher after the empiric CRM in Scenario 5. In Scenario 6, the logistic CRM has better PCS than the empiric CRM, with higher performance of our proposed methods.

Table 7: Proportions of selecting each dose regimen as the MTD-regimen over the 1000 trials in the six main toxicity scenarios after an empiric CRM. For each scenario, the marginal probabilities of DLT, CRS and DLT_o are defined, and the association between the CRS and DLT_o is represented by the average risk ratio (RR). Results are presented for the 3 joint approaches (DRtox_indep, DRtox_copula and DRtox_cond) and the empiric CRM. The proportions of correct selection (PCS) of the MTD-regimen are represented in bold.

Scenario	Set	RR	Method		S_1	S_2	S_3	S_4	S_5	S_6
1	A	1.85		p_T	0.10	0.14	0.18	0.30	0.45	0.60
				$p_T^{(1)}$	0.05	0.07	0.10	0.18	0.22	0.30
				$p_T^{(2)}$	0.06	0.08	0.10	0.17	0.34	0.50
			DRtox_indep	0	3	24	55	17	2	
			DRtox_copula	0	2	18	55	22	3	
			DRtox_cond	1	2	23	55	17	2	
			Empiric crm	0	4	23	50	20	2	
2	A	5.91		p_T	0.10	0.13	0.18	0.30	0.42	0.53
				$p_T^{(1)}$	0.07	0.10	0.13	0.22	0.28	0.36
				$p_T^{(2)}$	0.04	0.07	0.10	0.21	0.39	0.52
			DRtox_indep	0	4	28	50	14	3	
			DRtox_copula	0	2	20	54	19	5	
			DRtox_cond	0	4	24	52	16	3	
			Empiric crm	0	5	22	49	20	4	
3	A	1.81		p_T	0.11	0.15	0.18	0.30	0.45	0.59
				$p_T^{(1)}$	0.03	0.05	0.07	0.13	0.16	0.23
				$p_T^{(2)}$	0.08	0.11	0.13	0.22	0.39	0.54
			DRtox_indep	1	3	25	51	17	2	
			DRtox_copula	0	2	18	54	22	4	
			DRtox_cond	1	2	26	52	18	2	
			Empiric crm	1	5	27	48	18	2	
4	A	1.90		p_T	0.09	0.13	0.17	0.30	0.44	0.59
				$p_T^{(1)}$	0.07	0.10	0.13	0.23	0.29	0.37
				$p_T^{(2)}$	0.03	0.04	0.06	0.12	0.27	0.43
			DRtox_indep	0	3	25	55	15	1	
			DRtox_copula	0	2	19	57	19	3	
			DRtox_cond	0	4	24	56	15	1	
			Empiric crm	0	4	23	51	20	2	
5	A	1.97		p_T	0.03	0.04	0.05	0.11	0.17	0.03
				$p_T^{(1)}$	0.02	0.03	0.05	0.09	0.11	0.17
				$p_T^{(2)}$	0.00	0.01	0.01	0.02	0.08	0.18
			DRtox_indep	0	0	0	3	26	71	
			DRtox_copula	0	0	0	2	16	82	
			DRtox_cond	0	0	0	3	26	71	
			Empiric crm	0	0	0	2	20	77	
6	B	1.70		p_T	0.16	0.30	0.43	0.55	0.73	0.86
				$p_T^{(1)}$	0.09	0.17	0.25	0.33	0.38	0.48
				$p_T^{(2)}$	0.08	0.18	0.29	0.41	0.67	0.84
			DRtox_indep	22	61	16	0	0	0	
			DRtox_copula	14	62	23	2	0	0	
			DRtox_cond	21	60	17	1	0	0	
			Empiric crm	18	54	26	2	0	0	

Web Appendix C: Alternative set of dose regimens

In the main six toxicity scenarios, built either on Set A or B of the dose regimens that were inspired by the motivating trial, the CRS and DLT_o rarely occur at the same time. We then defined another set of dose regimens, Set C shown in Table 8, to increase the occurrence of both toxicities at the same time. In Set C, dose-escalation is slower than in Set A and B because the higher the steady-state dose is, the slower it is reached. We defined six additional toxicity scenarios on this new set that are similar to the main scenarios.

In Scenario 10 (similar to Scenario 1), the true MTD-regimen was S_4 that had a similar probability of CRS and DLT_o . The CRS and DLT_o were positively correlated with a average risk ratio of 1.89. In Scenario 11 (similar to Scenario 2), the association between the CRS and the DLT_o was increased to an average risk ratio of 6.37. In Scenarios 12 and 13 (similar to Scenarios 3 and 4), the true MTD-regimen remained S_4 but the proportion of each type of toxicity varied with a higher probability of DLT_o and CRS in Scenaris 12 and 13, respectively. Finally, the MTD-regimen changed to dose regimens S_6 and S_2 for Scenarios 14 and 15 (similar to Scenarios 5 and 6), respectively. The distribution of DLT, CRS and DLT_o per administration is illustrated in Figure 9 to compare the main scenarios built on Set A and B with the new scenarios built on Set C. For Scenarios 1, 5 and 6, most CRS occur at t_1 and t_4 while almost all DLT_o occur from t_4 as illustrated on the left part of the figure. However, for Scenarios 10, 14 and 15, both toxicities are more balanced throughout the drug administrations as illustrated on the right part on the figure even if CRS still tend to occur at the beginning while DLT_o occur at the end.

The proportion of selection of each dose regimen in these six new scenarios is shown in Table 9. In Scenarios 10, 11, 12 and 13, the PCS of the three joint approaches are similar to that of the CRM. All three joint approaches have a higher proportion of trials that recommend the under-dosing regimen, S_3 , as the MTD-regimen. In Scenario 14, the CRM has higher PCS that the joint approaches but it was already observed in the main scenarios. Finally, on Scenario 15, where the true MTD-regimen is S_2 , the three joint approaches outperform the CRM.

We represented in Figure 10 the estimated probabilities of DLT, CRS and DLT_o for the three joint approaches and the CRM for Scenario 10. We can observe that the marginal probability of CRS is slightly overestimated by the DRtox, while the probability of DLT_o (either the marginal probability of the conditional probability given no CRS) is estimated with a high variance for each joint approach. As a result, the probability of DLT at the MTD-regimen (S_4) and at the previous regimen (S_3) is also slightly overestimated by the joint approaches, which explains the higher proportion of trials that recommend S_3 as the MTD-regimen. In conclusion, in this new set of dose regimens, our joint modeling approaches do not improve the PCS compared to the CRM but they can still able to evaluate the probability of toxicity of new regimens, as illustrated in the main paper.

Table 8: Set A and Set C dose regimens used in the simulation study (in $\mu\text{g}/\text{kg}$).

		t_1	t_2	t_3	t_4	t_5	t_6	t_7
Set A	S_1	1	5	10	20	20	20	20
	S_2	1	5	10	25	25	25	25
	S_3	1	5	10	30	30	30	30
	S_4	1	5	10	45	45	45	45
	S_5	5	10	25	75	75	75	75
	S_6	10	25	50	100	100	100	100
Set C	S_1	1	1	1	1	1	1	1
	S_2	1	10	10	10	10	10	10
	S_3	1	10	30	30	30	30	30
	S_4	1	10	30	60	60	60	60
	S_5	1	10	30	60	100	100	100
	S_6	1	10	30	60	100	140	140




Figure 9: Mean number of DLT, CRS and DLT_o per trial for each of the seven administrations of the dose regimens observed in Scenarios 1, 10, 5, 14, 6, 15. Scenarios 1, 5 and 6 (Sets A and B) have similar probabilities of toxicities than Scenarios 10, 14, and 15 (Set C), respectively.

Table 9: Proportions of selecting each dose regimen as the MTD-regimen over the 1000 trials in the six toxicity scenarios defined on Set C. For each scenario, the marginal probabilities of DLT, CRS and DLT_o are defined, and the association between the CRS and DLT_o is represented by the average risk ratio (RR). Results are presented for the 3 joint approaches (DRtox_indep, DRtox_copula and DRtox_cond) and the logistic CRM. The proportions of correct selection (PCS) of the MTD-regimen are represented in bold.

Scenario	Set	RR	Method		S_1	S_2	S_3	S_4	S_5	S_6
10	C	1.89		p_T	0.03	0.09	0.18	0.30	0.45	0.57
				$p_T^{(1)}$	0.03	0.08	0.12	0.18	0.26	0.36
				$p_T^{(2)}$	0.00	0.01	0.07	0.17	0.31	0.41
			DRtox_indep	0	1	37	51	9	2	
			DRtox_copula	0	1	26	55	16	3	
			DRtox_cond	0	1	36	51	9	2	
			Logistic CRM	0	1	19	51	23	6	
11	C	6.37		p_T	0.04	0.10	0.18	0.30	0.45	0.58
				$p_T^{(1)}$	0.04	0.10	0.15	0.22	0.33	0.46
				$p_T^{(2)}$	0.00	0.01	0.06	0.22	0.41	0.55
			DRtox_indep	0	2	47	46	5	1	
			DRtox_copula	0	1	32	56	10	2	
			DRtox_cond	0	2	40	50	7	1	
			Logistic CRM	0	1	19	53	23	4	
12	C	5.84		p_T	0.02	0.09	0.18	0.30	0.42	0.52
				$p_T^{(1)}$	0.02	0.08	0.12	0.17	0.25	0.34
				$p_T^{(2)}$	0.00	0.02	0.11	0.27	0.41	0.51
			DRtox_indep	0	2	43	42	10	3	
			DRtox_copula	0	1	31	48	14	6	
			DRtox_cond	0	1	36	48	12	4	
			Logistic CRM	0	0	19	45	26	10	
13	C	1.94		p_T	0.04	0.11	0.19	0.30	0.45	0.61
				$p_T^{(1)}$	0.04	0.10	0.17	0.24	0.36	0.51
				$p_T^{(2)}$	0.00	0.01	0.04	0.10	0.19	0.28
			DRtox_indep	0	3	45	46	5	0	
			DRtox_copula	0	2	36	51	10	1	
			DRtox_cond	0	3	45	46	5	0	
			Logistic CRM	0	2	20	53	21	4	
14	C	1.99		p_T	0.01	0.05	0.08	0.11	0.18	0.30
				$p_T^{(1)}$	0.01	0.05	0.08	0.10	0.14	0.18
				$p_T^{(2)}$	0.00	0.00	0.00	0.00	0.04	0.17
			DRtox_indep	0	0	0	9	32	59	
			DRtox_copula	0	0	0	5	25	70	
			DRtox_cond	0	0	0	9	32	59	
			Logistic CRM	0	0	0	2	20	77	
15	C	1.82		p_T	0.16	0.30	0.44	0.57	0.75	0.88
				$p_T^{(1)}$	0.11	0.17	0.28	0.42	0.66	0.83
				$p_T^{(2)}$	0.06	0.18	0.26	0.33	0.41	0.47
			DRtox_indep	14	72	13	0	0	0	
			DRtox_copula	10	69	20	1	0	0	
			DRtox_cond	14	72	14	0	0	0	
			Logistic CRM	14	56	27	2	0	0	




Figure 10: Violin plots of the estimated probabilities of DLT, CRS and DLT_0 in Scenario 10 for 1000 simulated trials. All three joint approaches and the CRM estimate the probability of DLT in the first part of the figure, where the dashed line represents the toxicity target and the solid line represents the true DLT probabilities. Our three joint approaches estimate the probability of CRS with the logistic DRtox in the second part of the figure, where the solid line represents the true CRS probabilities. In the last part of the figure, both the DRtox_indep and DRtox_copula estimate the marginal probability of DLT_0 while the DRtox_cond estimates the conditional probability of DLT_0 given no CRS has occurred. The solid line represents the true marginal probabilities of DLT_0 while the dotted line represents the true conditional probabilities of DLT_0 given no CRS. In all the figure, horizontal lines on the density estimates represent the median and first and third quantiles of the distributions, and the plus sign represents the mean.