



HAL
open science

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

Emma Gerard, Sarah Zohar, Christelle Lorenzato, Moreno Ursino,
Marie-karelle Riviere

► **To cite this version:**

Emma Gerard, Sarah Zohar, Christelle Lorenzato, Moreno Ursino, Marie-karelle Riviere. Bayesian modeling of a bivariate toxicity outcome for early phase oncology trials evaluating dose regimens. *Statistics in Medicine*, 2021, 10.1002/sim.9113 . hal-03288022

HAL Id: hal-03288022

<https://hal.sorbonne-universite.fr/hal-03288022v1>

Submitted on 16 Jul 2021

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.

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'Ecole 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 (DLT_o) 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 $Y_{ij}^{(1)}$ be the binary CRS response and $Y_{ij}^{(2)}$ 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_{ij} is defined as $Y_{ij} = \max(Y_{ij}^{(1)}, Y_{ij}^{(2)})$. 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

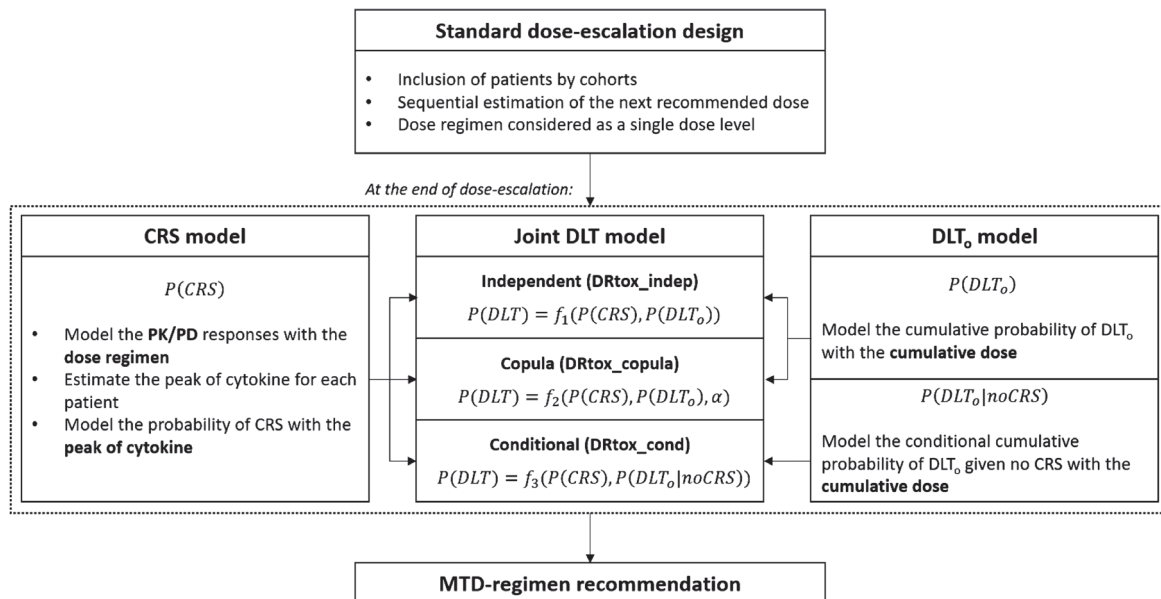


FIGURE 1 Illustration of the proposed modeling process of the bivariate toxicity endpoint at the end of the dose-escalation phase of the trial

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:

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

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

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

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 DLT_0 is only one type of the bivariate DLT, the number of observed DLT_0 is relatively small; thus we propose to reduce the number of parameters to two as follows:

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

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 DLT_0 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 DLT_0 . Additional details on prior elicitation can be found in Appendix A.

2.2.2 | Conditional DLT_0 modeling

The conditional DLT_0 model is very similar to the marginal DLT_0 model. We define the conditional cumulative probability of DLT_0 , given that no CRS has occurred, as follows:

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

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 DLT_0 model as

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

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

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

The posterior conditional probability of DLT_0 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 DLT_0 . 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 DLT_0 . 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 DLT₀ 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 DLT₀ 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 DLT₀, 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:

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

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:

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

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 DLT_o, we simulated the DLT_o 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 DLT_o than a patient without CRS. We therefore simulated the DLT_o from the conditional probability of DLT_o at each administration (1) given that no DLT_o 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 $d_{i,j}$ represents the dose given to patient i during the j^{th} administration and $D_{i,j} = \sum_{l=1}^j d_{il}$ 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 DLT_o 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 DLT_o during administration j (given that no DLT_o 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 DLT_o, $\lambda > 1$ represents positive association and $\lambda < 1$ represents negative association.

The marginal probability of DLT_o 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 DLT_o and the conditional probabilities of DLT_o given CRS and DLT_o 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 DLT_o and various associations between the CRS and the DLT_o. The latter association was measured for each scenario as the average risk ratio over the set of dose regimens of experiencing a DLT_o 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 DLT_o. The CRS and DLT_o 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 DLT_o than that of a patient not experiencing a CRS. In Scenario 2, the association between the CRS and the DLT_o 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 DLT_o 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 (DLT_0). For each scenario, the marginal probabilities of DLT, CRS, and DLT_0 of each dose regimen are represented as solid lines, while the conditional probabilities of DLT_0 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 DLT_0 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_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 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 DLT₀. For example, at cohort 5, three patients were included at dose regimen $S_4 = (1, 5, 10, 50, 50, 50, 50)$ $\mu\text{g/kg}$ 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 DLT₀ 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: $(1, 5, 10, 50)$ $\mu\text{g/kg}$. 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 .

DLT₀ model

1. The cumulative marginal (and conditional given no CRS) probabilities of DLT₀ 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 DLT₀ (and conditional posterior probabilities of DLT₀ given no CRS) are estimated. The marginal probabilities of DLT₀ (and the conditional probabilities of DLT₀ 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 DLT₀. 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 DLT₀ 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 DLT₀ 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 DLT₀, 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 DLT₀) in multiple dose administrations, that is, when dose regimens were administered, assuming that the CRS and the DLT₀ might be associated. We developed a modeling approach for the DLT₀ 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 1 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 1 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 $\pi_k = 1 - (1 - \pi_k^{(1)}) (1 - \pi_k^{(2)}) = 1 - (1 - \pi_k^{(1)})^2 = 1 - (1 - \pi_k^{(2)})^2$, 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 $\beta_{1,1} \sim \gamma(\alpha_1, \frac{\alpha_1}{\beta_{1,1}})$, 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:

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

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:

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

where T_{inf_j} 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 $t_{j+1} - t_j > T_{inf_j}$ 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:

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

The true marginal probability of CRS was computed by simulations using a Monte-Carlo method. The marginal probability of DLT_0 was computed as follows:

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